The JOREK non-linear extended MHD code and applications to large-scale instabilities and their control in magnetically confined fusion plasmas

M. Hoelzl[1,*], G.T.A. Huijsmans[2,3], S.J.P. Pamela[4], M. Bécoulet[2],
E. Nardon[2], F.J. Artola[1,5], B. Nkonga[6], C.V. Atanasiu[7], V. Bandaru[1],
A. Bhole[6], D. Bonfiglio[8], A. Cathey[1,9], O. Czarny[10], A. Dvornova[2],
T. Fehér[1], A. Fil[4], E. Franck[11], S. Futatani[12], M. Gruca[13], H. Guillard[14],
J.W. Haverkort[15], I. Holod[1,16], D. Hu[17], S.K. Kim[18], S.Q. Korving[3],
L. Kos[19], I. Krebs[20], L. Kripner[21], G. Latu[2], F. Liu[2], P. Merkel[1],
D. Meshcheriakov[1], V. Mitterauer[1,9], S. Mochalskyy[1], J.A. Morales[2],
R. Nies[1,22,23], N. Nikulsin[1,9], F. Orain[1], J. Pratt[24], R. Ramasamy[1,9,25],
P. Ramet[26], C. Reux[2], K. Särkimäki[1], N. Schwarz[1], P. Singh Verma[1],
S.F. Smith[4], C. Sommariva[27], E. Strumberger[1], D.C. van Vugt[3],
M. Verbeek[3], E. Westerhof[20], F. Wieschollek[1,9] and J. Zielinski[28]

[1] Max Planck Institute for Plasma Physics, Boltzmannstr. 2, 85748 Garching b. M., Germany
[2] CEA, IRFM, 13108 Saint-Paul-Lez-Durance, France
[3] Eindhoven University of Technology, PO Box 513, 5600 MB Eindhoven, Netherlands
[4] CCFE, Culham Science Centre, OX14 3DB, United Kingdom of Great Britain and Northern Ireland
[5] ITER Organization, 13067 St. Paul Lez Durance Cedex, France
[6] Université Côte d'Azur & Inria Sophia-Antipolis Méditerranée, CNRS, LJAD, Nice, France
[7] National Institute for Laser, Plasma and Radiation Physics, Atomistilor 409, PO BoxMG-36, 077125 Magurele-Bucharest, Romania
[8] Consorzio RFX-CNR, ENEA, INFN, Università di Padova, Acciaierie Venete SpA. I-35127 Padova, Italy
[9] Department of Physics, Technical University Munich, James-Franck-Str. 1, 85748 Garching b. M., Germany
[10] Framatome, 10 Rue Juliette Récamier, 69456 Lyon Cedex, France
[11] Inria Nancy Grand Est and IRMA Strasbourg, France
[12] Universitat Politècnica de Catalunya, Barcelona, Spain
[13] Institute of Plasma Physics and Laser Microfusion, PO Box 49, PL-00-908 Warsaw 49, Poland
[14] Inria Sophia-Antipolis Méditerranée & Université Côte d'Azur, Inria, CNRS, LJAD, Nice, France
[15] Delft University of Technology, Delft, Netherlands
[16] Max Planck Computing and Data Facility, Boltzmannstr. 2, 85748 Garching b. M., Germany
[17] School of Physics, Beihang University, Beijing, 100191, China
[18] Department of Mechanical and Aerospace Engineering, Princeton University, Princeton, NJ, United States of America
[19] LECAD Laboratory, Mech. Eng., University of Ljubljana, Aškerčeva 6, SI-1000 Ljubljana, Slovenia
[20] DIFFER—Dutch Institute for Fundamental Energy Research, De Zaale 20, 5612 AJ Eindhoven, Netherlands
[21] Institute of Plasma Physics, Prague, Czech Republic

*Author to whom any correspondence should be addressed.

[22] Department of Astrophysical Sciences, Princeton University, Princeton, NJ, 08543, United States of America

[23] Princeton Plasma Physics Laboratory, Princeton, NJ, 08540, United States of America

[24] Georgia State University, Department of Physics and Astronomy, Atlanta Georgia, 30303, United States of America

[25] Max-Planck/Princeton Research Center for Plasma Physics, Germany

[26] LaBRI, Inria Bordeaux Sud-Ouest, Université de Bordeaux, 351, cours de la Libération, 33405 Talence Cedex, France

[27] Ecole Polytechnique Fédérale de Lausanne (EPFL), Swiss Plasma Center (SPC), CH-1015 Lausanne, Switzerland

[28] University of Saskatchewan, S7N 5E2 Saskatoon, Canada

E-mail: matthias.hoelzl@ipp.mpg.de

CrossMark

Abstract

JOREK is a massively parallel fully implicit non-linear extended magneto-hydrodynamic (MHD) code for realistic tokamak X-point plasmas. It has become a widely used versatile simulation code for studying large-scale plasma instabilities and their control and is continuously developed in an international community with strong involvements in the European fusion research programme and ITER organization. This article gives a comprehensive overview of the physics models implemented, numerical methods applied for solving the equations and physics studies performed with the code. A dedicated section highlights some of the verification work done for the code. A hierarchy of different physics models is available including a free boundary and resistive wall extension and hybrid kinetic-fluid models. The code allows for flux-surface aligned iso-parametric finite element grids in single and double X-point plasmas which can be extended to the true physical walls and uses a robust fully implicit time stepping. Particular focus is laid on plasma edge and scrape-off layer (SOL) physics as well as disruption related phenomena. Among the key results obtained with JOREK regarding plasma edge and SOL, are deep insights into the dynamics of edge localized modes (ELMs), ELM cycles, and ELM control by resonant magnetic perturbations, pellet injection, as well as by vertical magnetic kicks. Also ELM free regimes, detachment physics, the generation and transport of impurities during an ELM, and electrostatic turbulence in the pedestal region are investigated. Regarding disruptions, the focus is on the dynamics of the thermal quench (TQ) and current quench triggered by massive gas injection and shattered pellet injection, runaway electron (RE) dynamics as well as the RE interaction with MHD modes, and vertical displacement events. Also the seeding and suppression of tearing modes (TMs), the dynamics of naturally occurring TQs triggered by locked modes, and radiative collapses are being studied.

Keywords: disruptions, edge localized modes, vertical displacement events, ELM control, disruption mitigation, MHD simulations, tokamak

(Some figures may appear in colour only in the online journal)

1. Introduction

The present article provides a comprehensive overview of the non-linear extended magneto-hydrodynamic (MHD) code JOREK, which is among the leading simulation codes worldwide for studying large scale plasma instabilities and their control in realistic divertor tokamaks. The article provides a detailed description of the physics models, numerical methods, and physics applications of the code.

In the existing literature, references [1, 2] already describe some aspects of the numerical methods and physics models of the JOREK code, which has been extended significantly since these articles were published. Reference [3] contains an overview of modelling activities worldwide regarding edge

localized modes (ELMs) and ELM control based on many different simulation codes, and references [4–6] provide a partial overview of JOREK activities regarding plasma edge and scrape off layer. The present article, in contrast, aims to give a comprehensive description of the code and its applications, with a particular focus on recent developments.

In the present section, we describe the motivation for the research activities (subsection 1.1) followed by a very brief review of (extended) magnetohydrodynamics (subsection 1.2) and some words on the historic development of JOREK (subsection 1.3).

The rest of the article is organized as follows: section 2 provides a detailed overview of the physics models available in JOREK and section 3 describes the numerical methods employed for solving the equations. Selected tests performed for code verification are shown in section 4. After this 'technical' part, a detailed picture is drawn of the physics studies and validation activities performed in particular in the fields of plasma edge and scrape off layer physics (section 5) as well as disruption physics (section 6). Further code applications are described in section 7. Each section contains a brief outlook towards further plans and developments. Finally, a concise summary is provided in section 8.

The support received from many entities and useful discussions with various scientists are acknowledged. Additional details on the coordinate systems, finite element basis, normalization of quantities, and time stepping scheme are provided in appendices A–C.

1.1. Motivation and challenges

Among the obstacles, which need to be overcome on the path towards a magnetic confinement fusion power plant, large scale plasma instabilities may well be the most critical one. A plasma configuration suitable for harvesting energy needs to have good confinement properties, however, a reliable control[29] of plasma instabilities is equally important. Robust predictions of the properties of such instabilities and of effective control methods are urgently needed

  (a) To provide input to the ITER design, where it can still be influenced [e.g., the disruption mitigation system (DMS)],
  (b) To prepare a robust, efficient, and successful exploitation of ITER across all phases of the planned operation, and
  (c) To answer critical questions regarding the design of a successful DEMO reactor.

Revealing the underlying physics processes of plasma instabilities and developing control mechanisms constitutes a major challenge for experiments, theory, and modelling. While the suitability of control techniques for present devices may be tested in a straight forward manner experimentally, their applicability to future machines, with plasma parameters very different both quantitatively (e.g., Lundquist number) and qualitatively (e.g., large amount of fusion-born fast particles) from present machines, needs to be ensured by developing

truly predictive capabilities. In such a holistic approach based on fundamental plasma theory, experimental studies across devices, and numerical simulations of the plasma dynamics, the computational models play a key role. Simulation codes can provide the capability to predict the relevant processes in future devices after being carefully validated against theory predictions and experiments first. Activities with the JOREK code ultimately aim at reaching that goal.

Key challenges in this respect are the immense scale separations in both time and space of the involved processes, the intrinsic highly non-linear multiphysics nature of the problem, and the complicated magnetic topology of divertor plasmas. MHD models have become a very robust and reliable framework for describing large-scale plasma instabilities. And via numerous extensions beyond the classical MHD, more and more effects can be captured accurately in the simulations. Worldwide, a number of specialized simulation codes for calculating non-linear MHD dynamics in magnetically confined tokamak and stellarator plasmas have been developed in the past years and decades including BOUT++ [7], JOREK (this article and references [1, 2, 8]), MEGA [9], M3D [10], M3D-C^1 [11–13], NIMROD [14, 15], and XTOR [16, 17] (listed alphabetically, not a complete list).

Besides the challenges imposed by the multi-scale nature already mentioned, in particular the large number of different physical effects, which need to be treated consistently and which are mutually interacting in a highly non-linear way, requires simulation codes that can capture this rich multiphysics behaviour in a reliable way. In a typical mitigated disruption scenario, for instance, the dynamics of magnetic islands, the ablation of (shattered) pellets, the reconnection of the plasma leading to a stochastic state, the fast losses of thermal energy along magnetic field lines, the radiative losses by partly ionized impurities, the generation and transport of runaway electrons (REs), the interaction of REs with the MHD modes and the electromagnetic interaction of the plasma with conducting structures in the device may all play an important role simultaneously. Developing the capability to describe the non-linear interaction of all these processes is necessary for unravelling the complete physics picture and becoming truly predictive regarding the dynamics in future machines. At the same time, simpler models are needed to allow faster access to larger parameter studies. JOREK is an advanced simulation framework for studying large-scale instabilities in magnetized plasmas. It offers such a hierarchy from simple and fast to very complex and computationally demanding models.

1.2. Extended magnetohydrodynamics (MHD)

This article does not give a complete overview of magnetohydrodynamics (MHD) and its computational treatment. We mention only key features in this section, which are directly relevant as context for this article. For literature on MHD, in particular the references [18–24] are recommended.

MHDs developed first by H. Alfvén in 1942 [25] describes a magnetized plasma as an electrically conducting fluid. In the ideal MHD model, the plasma is assumed to be perfectly

---

[29] The term 'control' is in this article is meant to include both avoidance and control strategies.

conducting. Ideal MHD can describe certain stability limits in tokamak plasmas well (e.g., see references [26–29] for type-I ELMs). However, 3D non-linear simulations need to be based on resistive extended MHD models, which include anisotropic heat conduction, plasma resistivity, diamagnetic flows, finite Larmor radius effects, neoclassical physics, source/sink terms, two-fluid effects, neutrals, impurities, sheath boundary conditions, and many more effects depending on the addressed problem. A certain class of models includes also electron inertia effects [30].

Tokamak plasmas are typically in approximate force balance $\nabla p \approx \mathbf{j} \times \mathbf{B}$, where p denotes (the isotropic component of) the pressure, $\mathbf{j}$ the plasma current vector, and $\mathbf{B}$ the magnetic field vector. The stability of this equilibrium state determines whether the plasma will remain in this equilibrium state or is prone to instabilities. This is traditionally studied by linearizing the equations and analyzing the eigenvalue spectrum of the system along with the associated eigenvectors. However, linear growth rates may be affected dramatically by background flows and non-ideal plasma effects, which are not always accounted for in linear codes. Also, non-linear dynamics cannot be predicted from the linear stability analysis in general, and linearly stable eigenmodes might become non-linearly unstable at sufficiently large 'seed perturbation' amplitudes [e.g., neoclassical tearing modes (NTM)]. As a result, predicting the full consequences of plasma instabilities is only possible by employing advanced non-linear models. Solving such models in realistic geometries typically is only possible numerically. MHD involves very different time scales: the Alfvèn time $\tau_{\mathrm{A}} = a\sqrt{\mu_0 m_{\mathrm{i}} n_{\mathrm{i}}}/B$ is about 0.3 μs for ITER like parameters, where a denotes the minor radius of the plasma, μ_0 the vacuum permeability, m_{i} the ion mass, and n_{i} the ion density. On the other hand, the resistive time scale $\tau_{\mathrm{R}} = \mu_0 a^2/\eta$, where η denotes the plasma resistivity, is $\gg 1$ s for ITER like parameters. Plasma instabilities typically develop on mixed time scales of tens of μs to tens of ms. The resistive time scale of the ITER vacuum vessel is around $\tau_{\mathrm{w}} = 500$ ms slowing down some instabilities to that time scale (e.g. axisymmetric resistive wall modes). Consequently, the relevant time scales for large scale instabilities are two to six orders of magnitude longer than the Alfvén time. The frequencies of fast magneto-acoustic waves propagating in the plane orthogonal to the magnetic field are typically even two to three orders of magnitude larger than the Alfvén frequency, thus constituting the most challenging time scale in the system. The so-called reduced MHD model, described in section 2.3.1, eliminates the fast waves from the model to facilitate its numerical solution.

In spatial dimensions, a similarly challenging splitting of scales can be observed. While the size of the whole system typically is in the range of several metres (minor radius of 2 m in ITER), the resistive skin depth is given by $\sqrt{2\eta/(\mu_0\omega)}$ at a given frequency ω, which can easily drop into the mm or even sub-mm range at the low resistivity of large fusion devices (which decreases strongly with temperature)—a separation by four orders of magnitude. The strong increase of this scale separation towards larger (and at the same time hotter) fusion devices is a particular challenge for the modelling.

Anisotropic heat conduction is another particularly challenging physics aspect to be dealt with in MHD simulations. While the transport coefficients across field lines determined by neoclassical or turbulent processes typically are in the range of 1 $\mathrm{m}^2\,\mathrm{s}^{-1}$, the heat transport along field lines by electrons can reach values of $10^{10}\,\mathrm{m}^2\,\mathrm{s}^{-1}$ in hot plasmas [31, 32]. Avoiding overly restrictive time scales, numerical instabilities, or a pollution of cross-field transport by errors in the parallel transport is a significant challenge for the numerical treatment.

Magnetohydrodynamics is strictly valid only when the plasma is sufficiently collisional, and many important kinetic effects are not reflected by the MHD equations. However, a large number of corrections (e.g., effective parallel heat diffusion coefficients [32]) and extra terms (e.g., two-fluid effects, or a consistent evolution of the bootstrap current [33]) allow to apply MHD outside its original boundaries. In many cases, the full MHD equations can be further simplified to eliminate the fast magneto-sonic waves from the system, reducing the separation of time scales. A significant number of reduced MHD models with different levels of approximation exist (e.g., references [34, 35]), which lower the number of physical variables in the system. JOREK presently has several different reduced MHD models (the one described in section 2.3 with and without parallel velocity; a reduced MHD model suitable for stellarator applications is in development) and a full MHD model (section 2.12) implemented for tokamak configurations along with numerous physics extensions.

1.3. Historic development of the JOREK code

The development of a first version 'JOREK 1' was started by G.T.A. Huysmans in 2002 at CEA/IRFM and is described in reference [36]. Applications of the JOREK 1 code include the current hole problem, the stability of external kink modes in X-point plasmas [36], the first nonlinear ELM simulations [1] and the application of resonant magnetic perturbation (RMP) fields [37]. The JOREK 1 code was based on so-called generalised, h–p refinable, finite elements [38]. However, in practice, the p refinement, i.e., adapting the order of the finite elements was never used. Therefore it was decided to change the finite elements to cubic Bezier finite elements, an extension to the iso-parametric bicubic Hermite elements which are successfully applied in the HELENA equilibrium code [39]. The code 'JOREK 2', which has been developed since 2006, is first described in the references [2, 40, 41] and has successively evolved into the presently existing JOREK code, which is described in the article at hand. As major changes in version 2, an iterative solver, and a G^1 continuous finite element formulation had been implemented. G^1 continuity refers to both values and real-space gradients being continuous throughout the computational domain but without continuity in the gradients in the local finite element coordinates. The present JOREK code is being further developed continuously regarding physics models, numerical methods, and applications as shown in this article. The JOREK website [42] contains some regularly updated information. The article at hands intends to give a complete overview of the code including references
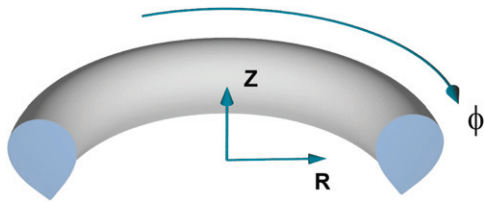
Figure 1. The base cylindrical coordinate system used in JOREK. Refer to appendix A for details.

to all original publications which go more into detail than possible here.

2. Physics models

This section describes the physics models and corresponding extensions available in JOREK. Before turning towards these models, the coordinate systems used in JOREK are introduced briefly.

2.1. Coordinate systems

The base cylindrical coordinate system (R, Z, ϕ) is given by $x = R \cos \phi$, $y = -R \sin \phi$, $z = Z$, where (x, y, z) denotes Cartesian coordinates (figure 1). Thus, ϕ is oriented clockwise if viewed from the top. According to the definitions in reference [43], the JOREK conventions correspond to a COCOS number of 8. To describe the Bezier elements, the coordinates R and Z are expanded in the same Bezier basis functions, that are also used for the expansion of the physics variables ('isoparametric'). This introduces a local (s, t, ϕ) coordinate system inside each grid element. See section 3.1 and appendix A for more details on the discretization.

2.2. Grad–Shafranov solver

JOREK has a built-in Grad–Shafranov equilibrium solver which uses the same finite element grid and representation of the variables used in the nonlinear time evolution. This guarantees that the discrete initial state used in MHD equations accurately satisfies the initial equilibrium force balance, avoiding any initial discontinuous behaviour. JOREK can solve both fixed boundary equilibria and, through the coupling to the STARWALL code, free boundary equilibria (see section 2.9).

The solver requires the profiles of pressure (provided by temperature and density separately, since they are needed for the initial conditions) and FF'. These profiles are provided as functions of the normalized poloidal flux $\Psi_N = (\Psi - \Psi_{\text{axis}})/(\Psi_{\text{bnd}} - \Psi_{\text{axis}})$, either via a simple analytical function, or via a numerical representation. Here Ψ_{axis} and Ψ_{bnd} denote the values of the poloidal magnetic flux at the magnetic axis and on the boundary of the plasma domain, respectively. In addition, the poloidal flux Ψ on the boundary of the computational domain needs to be specified by a numerical list of (R, Z, Ψ) points (or by coefficients for analytical moments for simpler cases). This input can be extracted, for instance, from 'geqdsk' files or from equilibria created with the CLISTE

code. Starting from an initial guess, the equilibrium is determined iteratively by Picard or Newton iterations to a specified accuracy. After solving the GS equation on the initial finite element grid, the solution is typically used to create a new grid aligned to the equilibrium flux surfaces. The GS equation is solved a second time on this new grid, providing the accurate initial conditions for the time evolution part.

After the equilibrium calculation, all physical variables are initialized consistently to it: the poloidal flux Ψ is directly taken from the equilibrium solution; the toroidal current density is calculated directly from Ψ via the current definition equation; density and temperature are initialized according to the specified profiles. All velocity related quantities (velocity stream function, vorticity, and parallel velocity) are initialized to be zero unless background rotation profiles are prescribed. In case of sheath boundary conditions, the parallel velocity is initialized to the ion sound speed at divertor targets[30].

2.3. Base MHD model

The MHD model is formulated as a set of normalized equations for the evolution of the magnetic potential (**A**), mean velocity (**V**), total density (ρ) and total pressure (p). The equations are normalized with respect to the central mass density ρ_0 and the vacuum permeability μ_0 such that μ_0 does not appear explicitly. Length scales are not normalized, while the time is normalized by a factor[31] $\tau_{\text{norm}} = \sqrt{\mu_0 \rho_0}$ which is typically close to the Alfvén time $\tau_A = a\sqrt{\mu_0 \rho_0}/B_0$. The total pressure and mass density are normalized by μ_0 and ρ_0 respectively. Details on the normalization of further quantities are given in appendix B. The normalized equations are written as

$$\frac{\partial \mathbf{A}}{\partial t} = -\mathbf{E} - \nabla \Phi, \tag{1}$$

$$\rho \frac{\partial \mathbf{V}}{\partial t} = -\rho \mathbf{V} \cdot \nabla \mathbf{V} - \nabla p + \mathbf{J} \times \mathbf{B} + \nabla \cdot \underline{\tau} + \mathbf{S_V}, \tag{2}$$

$$\frac{\partial \rho}{\partial t} = -\nabla \cdot (\rho \mathbf{V}) + \nabla \cdot (\underline{D}\nabla \rho) + S_\rho, \tag{3}$$

$$\frac{\partial p}{\partial t} = -\mathbf{V} \cdot \nabla p - \gamma p \nabla \cdot \mathbf{V} + \nabla \cdot (\underline{\kappa}\nabla T)$$
$$+ (\gamma - 1)\underline{\tau} : \nabla \mathbf{V} + S_p \tag{4}$$

The total pressure is defined by the ideal gas law $p = p_{\text{i}} + p_{\text{e}} = \rho T$ (μ_0 disappears due to normalization). This pressure is the sum of electron (p_{e}) and ion (p_{i}) pressures. Electron and ion pressures are then assumed in some of the models to be half of the total pressure, e.g., where the diamagnetic terms are calculated. The model extension described in section 2.5 allows to separately evolve variables for the electron and ion temperatures instead. The magnetic field vector, **B**, and the current

---

[30] Simulations with sheath boundary conditions typically need to be run axisymmetrically for a short while, such that the parallel flows in the SOL can establish a steady state. Non-axisymmetric Fourier modes are added then, once SOL flows have equilibrated.

[31] In case of ITER with minor radius $a \approx 2$ m and magnetic field amplitude on axis $B_0 \approx 5$ T, $\tau_{\text{norm}} \approx 2.5\tau_A$. Note that $\sqrt{\mu_0 \rho_0}$ does not have the dimension of a time and it is therefore more exact to say that the numerical value of τ_{norm} is $\sqrt{\mu_0 \rho_0}$.

vector, $\mathbf{J}$, are defined as:

$$\mathbf{B} = F\nabla\phi + \nabla \times \mathbf{A} \quad \text{and} \quad \mathbf{J} = \nabla \times \mathbf{B}. \tag{5}$$

Here, the toroidal flux function $F(\psi) \equiv RB_\phi$ is not essential for the model but is added for numerical reasons. $F(\psi)$ is constant in time and typically taken from the initial Grad–Shafranov equilibrium such that the initial vector potential in the poloidal plane is zero. This does not constrain $\mathbf{B}$ in any way since the magnetic vector potential $\mathbf{A}$ takes into account the (arbitrarily large) time evolution and perturbations from the equilibrium.

Ohm's law includes resistivity and drift-ordered (diamagnetic) terms [44]:

$$\mathbf{E} = -\mathbf{V} \times \mathbf{B} + \eta(\mathbf{J} - \mathbf{J}_\star) + F_0\frac{\delta^*}{\rho}(\nabla_\perp p_\mathrm{i} - \nabla_\parallel p_\mathrm{e}) \tag{6}$$

The resistivity η with Spitzer temperature dependence and the diamagnetic coefficient δ^* are given by

$$\eta = \eta_0 \cdot (T/T_0)^{-3/2} \tag{7}$$

$$\delta^* = m_\mathrm{ion}/(eF_0\sqrt{\mu_0\rho_0}) \tag{8}$$

with the constant parameter η_0 and initial plasma core temperature (T_0). m_ion is the ion mass and e the elementary charge. The constant $F_0 = R_0B_{\phi 0}$ is defined as the major radius at the geometric centre times the vacuum toroidal field. This constant appears due to the definition of the δ^*, for consistency with the reduced MHD model described below. A more accurate modelling of the diamagnetic effect is possible with the two pressures extension described in section 2.5. The term $\mathbf{J}_\star = j_S\mathbf{e}_\phi$ denotes a toroidal current source term. It can be used to preserve the original current profile j_0 approximately throughout the simulation, if one chooses $j_S = j_0$.[32] The current source term is also used to model a consistently evolving bootstrap current [45]. In that case, the initial current profile needs to include the initial bootstrap current correctly and the current source term takes the form: $j_S = j_0 + j_B - j_{B,0}$, where $j_{B,0}$ denotes the initial bootstrap current and j_B is the bootstrap current corresponding to the self-consistent profiles during time evolution. For the calculation of the bootstrap current, the expressions of reference [33] are used.

In the MHD model shown in equations (1)–(4), the gauge still needs to be defined. In the JOREK full MHD model (section 2.12), the Weyl gauge, $\Phi = 0$, is used. That implies that the toroidal component of the magnetic vector potential changes with time even in steady state.

The tokamak plasma evolves in a low collisionality regime and the associated viscous stress tensor ($\underline{\tau}$) is decomposed into three main parts

$$\underline{\tau} \simeq \underline{\tau}_f + \underline{\tau}_\mathrm{neo} + \underline{\tau}_{gv} \tag{9}$$

These components model the Newtonian-fluid type, neoclassical and gyro-viscous effects respectively. The separation in the different physical effects contained in the stress tensor

does not imply a possible double counting but rather a representation on different time scales and directions. Following [46], the parallel viscous stress is included as a Braginskii viscous stress acting on a fast timescale and a residual neoclassical stress on a slower collisional time scale. The Newtonian stress tensor ($\underline{\tau}_f$) is decomposed into the parallel and the perpendicular directions to the magnetic field and the associated coefficients of viscosity are $\nu_\parallel$ and $\nu_\perp$. According to the Chew–Goldberger–Low formulation [47], the parallel stress tensor $\underline{\tau}_{f\parallel}$ for arbitrary collisionality in a magnetized plasma is written as [20]:

$$\underline{\tau}_{f\parallel} = 3\nu_\parallel \left(\mathbf{b} \otimes \mathbf{b} - \frac{1}{3}\underline{I}\right) \otimes \left(\mathbf{b} \otimes \mathbf{b} - \frac{1}{3}\underline{I}\right)\nabla\mathbf{V}$$

The coefficient $\nu_\parallel$ is modelled as a spatial constant but such a dependency can be changed easily. The explicit formulation of the perpendicular tensor $\underline{\tau}_{f\perp}$ can be found in reference [20]. The associated coefficient $\nu_\perp$ is typically chosen to have the same temperature dependence as η in order to keep the magnetic Prandtl number spatially nearly constant (except for the weaker density dependency).

$$\nu_\perp = \nu_{\perp 0} \cdot (T/T_0)^{-3/2} \tag{10}$$

where $\nu_{\perp 0}$ is a constant parameter and T_0 is the initial plasma core temperature. The neoclassical viscous tensor ($\underline{\tau}_\mathrm{neo}$) is determined by a heuristic formulation [48]

$$\nabla \cdot \underline{\tau}_\mathrm{neo} = \rho\nu_\mathrm{neo}\frac{\|\mathbf{B}\|^2}{\|\mathbf{B}_\theta\|^2}(\mathbf{b}_\theta \otimes \mathbf{b}_\theta)(\mathbf{V} - \mathbf{V}_\mathrm{neo}) \quad \text{where}$$

$$\mathbf{b}_\theta = \frac{\mathbf{B}_\theta}{\|\mathbf{B}_\theta\|} \tag{11}$$

with $\mathbf{B}_\theta = \mathbf{B} - (\mathbf{B} \cdot \mathbf{e}_\phi)\mathbf{e}_\phi$ the poloidal magnetic field. The neoclassical coefficient ν_neo and velocity $\mathbf{V}_\mathrm{neo}$ are given functions of the temperature and the magnetic field (see references [49, 50]). In magnetized plasmas it is usual to assume gyro-viscous cancelation [20, 51] caused by the finite Larmor-radius effect. Therefore, to enforce gyro-viscous cancelation, the gyro-viscous stress tensor $\underline{\tau}_{gv}$ is modelled as

$$\nabla \cdot \underline{\tau}_{gv} = \rho \left(\frac{\partial\mathbf{v}_\mathrm{i}^*}{\partial t} + \mathbf{V} \cdot \nabla\mathbf{v}_\mathrm{i}^* + \mathbf{v}_\mathrm{i}^* \cdot \nabla\mathbf{v}_\parallel\right) \tag{12}$$

where $\mathbf{v}_\mathrm{dia,i}$ is the ion diamagnetic drift velocity defined in equation (26) and $\mathbf{v}_\parallel$ the parallel velocity. The terms involving the ion-diamagnetic heat flux (and the associated cancelation with the density convection due to the ion diamagnetic drift [52]) have not been implemented to avoid the possible destabilisation of ion temperature gradient (ITG) modes. These terms are essential to study the interaction of MHD instabilities with underlying ITG turbulence but this is left for future applications.

The heat diffusion tensor $\underline{\kappa}$ is decomposed parallel and perpendicular to the magnetic field

$$\underline{\kappa} = \kappa_\parallel\mathbf{b} \otimes \mathbf{b} + \kappa_\perp(\underline{I} - \mathbf{b} \otimes \mathbf{b}) \tag{13}$$

Note that the factor $(\gamma - 1)$ in the heat diffusion terms is absorbed in the coefficients $\kappa_\parallel$ and $\kappa_\perp$. Here, γ is the ratio of

---

[32] For a more consistent treatment, a loop voltage can also be applied at the computational boundary.

specific heats (usually $\gamma = 5/3$). The vector $\mathbf{b} = \mathbf{B}/B$ denotes the unit vector in the direction of the magnetic field. Radial profiles of $\kappa_\perp$ are usually specified in an ad hoc manner to mimic the background transport that cannot be captured with the present model. For instance, low values are set in the pedestal region to model the transport barrier. The parallel heat diffusion coefficient $\kappa_\parallel$ is implemented with Spitzer–Härm [31] temperature dependency according to

$$\kappa_\parallel = \kappa_{\parallel,0} \cdot (T/T_0)^{5/2}, \tag{14}$$

where the central value $\kappa_{\parallel,0}$ is calculated according to the Spitzer–Härm formula. An optional parameter $\kappa_{\parallel,\,\text{max}}$ can be specified to account for the heat flux limit [32] in a simplified way by ensuring that the parallel heat conductivity cannot exceed this maximum value. Realistic anisotropies even beyond $\kappa_\parallel/\kappa_\perp \approx 10^{10}$ can be handled without producing large spurious perpendicular transport provided a grid is used that is aligned to the equilibrium flux surfaces (see section 4.2). The particle diffusion tensor $\underline{D}$ has an analogous form to expression (13) although the parallel component ($D_\parallel$) is usually not used as the parallel particle transport is dominated by convection. The profile of $D_\perp$ is also specified by ad hoc profiles reflecting underlying small-scale turbulence that is not included in the MHD model.

The source term in the momentum equation contains the contribution of the diffusion and the source of density, as well as specific source $\mathbf{S_m}$ of momentum

$$\mathbf{S_V} = \mathbf{S_m} - \left(\nabla \cdot (\underline{D}\nabla\rho) + S_\rho \right) \mathbf{V}, \tag{15}$$

The source terms in the pressure equation contain the Ohmic heating term, thermal energy source and particles source effects.

$$S_p = (\gamma - 1) \Bigg((\mathbf{E} + \mathbf{V} \times \mathbf{B}) \cdot \mathbf{J} + S_\text{E} - \mathbf{V} \cdot \mathbf{S_m}$$
$$+ \frac{\mathbf{V} \cdot \mathbf{V}}{2} (S_\rho + \nabla \cdot (\underline{D}\nabla\rho)) \Bigg), \tag{16}$$

where S_E is the thermal energy source and S_ρ is the particle source. Sources are typically specified as radial profiles.

Given equations (1)–(4), a proper mathematical treatment of this system should specify the functional spaces where the solutions are sought for. Moreover, since JOREK uses a finite element method, a weak form of the equation is preferred which implies to define basis and test functions. In the following, for brevity, the 'dV' in all volume integrals is omitted. Thus let $\mathcal{V_A}, \mathcal{V_V}, \mathcal{V_\rho}, \mathcal{V_p}$ and $\mathcal{V_A^*}, \mathcal{V_V^*}, \mathcal{V_\rho^*}, \mathcal{V_p^*}$ be the chosen function spaces for the basis and test functions respectively, a weak form of the MHD problem will be reformulated as: find $(\mathbf{A}, \mathbf{V}, \rho, T)$ in $\mathcal{V_A} \times \mathcal{V_V} \times \mathcal{V_\rho} \times \mathcal{V_p}$ such that, for any test functions $(\mathbf{A^\star}, \mathbf{V^\star}, \rho^\star, p^\star)$ in $\mathcal{V_A^*} \times \mathcal{V_V^*} \times \mathcal{V_\rho^*} \times \mathcal{V_p^*}$, we have:

$$\int \frac{\partial \mathbf{A}}{\partial t} \cdot \mathbf{A^\star} = - \int \mathbf{E} \cdot \mathbf{A^\star}, \tag{17}$$

$$\int \rho \frac{\partial \mathbf{V}}{\partial t} \cdot \mathbf{V^\star} = - \int (\rho \mathbf{V} \cdot \nabla\mathbf{V} + \nabla p - \mathbf{J})$$
$$\times \mathbf{B} - \nabla \cdot \underline{\tau} - \mathbf{S_V}) \cdot \mathbf{V^\star}, \tag{18}$$

$$\int \frac{\partial \rho}{\partial t} \rho^\star = - \int \left(\nabla \cdot (\rho \mathbf{V}) - \nabla \cdot (\underline{D}\nabla\rho) - S_\rho \right) \rho^\star, \tag{19}$$

$$\int \frac{\partial p}{\partial t} p^\star = - \int (\mathbf{V} \cdot \nabla p + \gamma p \nabla \cdot \mathbf{V} - \nabla \cdot (\underline{\kappa}\nabla T)$$
$$- (\gamma - 1)\underline{\tau} : \nabla\mathbf{V} - S_p) \, p^\star \tag{20}$$

where $p = T\rho$. The identity $\mathbf{J} \times \mathbf{B} = -\nabla \left(\frac{\mathbf{B} \cdot \mathbf{B}}{2} \right) + \nabla \cdot (\mathbf{B} \otimes \mathbf{B})$ and integration by parts is used to avoid computation of second order derivatives. Equations (17) and (18) are vector equations. For numerical purpose, each of them must be transformed into three scalar equations by projecting the vectors onto some basis.

Following the representation of $\mathbf{A}$ and $\mathbf{V}$ in the basis $(\mathbf{e}_R, \mathbf{e}_Z, \mathbf{e}_\varphi)$, and $(\mathbf{e}_R, \mathbf{e}_Z, \mathbf{B})$ respectively, the basis for the vector-test function $\mathbf{V^\star}$ and $\mathbf{A^\star}$ in the weak formulation (17)–(20) is chosen as:

$$\mathbf{A} = A_R\mathbf{e}_R + A_Z\mathbf{e}_Z + \frac{1}{R}A_3\mathbf{e}_\varphi \tag{21}$$

$$\mathbf{V} = V_R\mathbf{e}_R + V_Z\mathbf{e}_Z + V_\parallel\mathbf{B} \tag{22}$$

$$\mathbf{A^\star} = a^* \left(\mathbf{e}_R, \mathbf{e}_Z, \mathbf{e}_\varphi \right) \tag{23}$$

$$\mathbf{V^\star} = v^* \left(\mathbf{e}_R, \mathbf{e}_Z, \mathbf{B} \right) \tag{24}$$

where a^* and v^* represent the scalar test functions as defined in section 3.1.1. This choice of projection in the parallel direction ensures on the discrete level that the Lorentz force $\mathbf{V}^* \cdot (\mathbf{J} \times \mathbf{B})$ is exactly vanishing.

2.3.1. Reduced MHD. In order to reduce computational requirements, one often employs reduced MHD models, which eliminate fast magnetosonic waves while retaining the relevant physics [12, 34, 35, 53]. The removal of fast magnetosonic waves, the fastest waves in the system, allows one to use larger time steps due to relaxing the CFL condition. Even when implicit time integration methods are used, and the CFL condition is no longer a hard limit, using time steps that are large compared to the shortest time scale can lead to poor accuracy [12, 54]. In addition, reduced MHD has less unknowns compared to full MHD, which decreases the computational costs and memory requirements for simulations.

Reduced MHD, as first introduced by Greene and Johnson [55], and later developed by Kadomtsev, Pogutse and Strauss [56, 57], relied on ordering in a small parameter, often taken to be the inverse aspect ratio. The ordering itself is a system of several approximations and assumptions involving the ordering parameter that allows one to determine the relative order (in terms of the ordering parameter) of any quantity with respect to any other quantity of the same dimension. In this context, terms corresponding to fast magnetosonic waves have a higher order than the terms that one wants to keep, allowing the fast wave terms to be dropped. Naturally, there are many choices one can make in the ordering assumptions, depending on which physical effects one wants to keep, all of which result in different reduced equations [35, 54, 57–59]. The ideas of reduced MHD have also found use in astrophysics, where toroidal geometry cannot be assumed, and thus the inverse aspect ratio cannot be used as an ordering parameter [60].

Starting in the 1980s, a new ansatz-based approach was introduced by Park *et al* [61], where an ansatz form that eliminates fast magnetosonic waves is used for the velocity and terms of all orders are kept (eliminating partially the fluid compression). Their reduced model corresponds to ideal MHD in the incompressible limit and was used to resolve internal kink modes in a cylindrical geometry, something that ordering-based reduced MHD could not do. Izzo *et al* used a similar ansatz in their study [62]. Later papers also adopt an ansatz for the magnetic field that eliminates field compression [63, 64]. The ansatz approach allows one to make less assumptions and keep more physical effects, while generally resulting in more complicated equations than the ordering approach. Thus, while keeping more physics, the various terms in the equations of ansatz-based reduced MHD are harder to interpret due to their complexity. In addition, without an ordering parameter, error estimation becomes much more difficult. The ansatz method allows to conserve energy exactly on the equation level. Some of the ordering-based models also conserve energy, however it is then often necessary to keep selected higher order terms to ensure that. The ansatz method thus is an alternative way of eliminating fast magnetosonic waves, which makes energy conservation easier. At the same time, in comparison to models which use $\epsilon = a/R$ for the ordering, the ansatz based reduced MHD model is still fairly accurate in the spherical tokamak limit as shown in reference [65] by comparison to full MHD.

The reduced MHD model used in JOREK is derived following the ansatz-based approach. In this approach, instead of the whole functional spaces used in full MHD, the variables are constrained to lie in a subset of these spaces and the equations are established by a Galerkin truncation. Another way to present this procedure is to say that an ansatz is postulated for some variables. The ansatz considered here assumes that the time dependent part of the magnetic potential is dominated by the toroidal component. The ansatz for the magnetic field is deduced by approximating B_ϕ as the vacuum F_0/R toroidal field.

$$\mathbf{B} = \frac{F_0}{R}\mathbf{e}_\phi + \frac{1}{R}\nabla\psi \times \mathbf{e}_\phi \Longrightarrow \mathbf{A} = \psi\nabla\phi \qquad (25)$$

where F_0 is constant in space as well as time and $\mathbf{e}_\phi$ is the normalized toroidal basis vector. In the weak formulation (17), this corresponds to defining $\mathcal{V}_\mathbf{A} = \{\mathbf{A} : \exists \psi \in H^2$ s.t. $\mathbf{A} = \psi\nabla\phi\}$. The ansatz (25) implies that the velocity cannot be arbitrary. Indeed, taking the cross product of equation (1) with $\nabla\phi$, after substituting equations (6) and (25) for $\mathbf{E}$ and $\mathbf{A}$ and neglecting resistivity and the poloidal component of $\mathbf{B}$, we obtain:

$$\mathbf{v}_P = \underbrace{-R\nabla u \times \mathbf{e}_\phi}_{\equiv \mathbf{v}_E} - \underbrace{(\delta^* R/\rho)\nabla p_i \times \mathbf{e}_\phi}_{\equiv \mathbf{v}_{\mathrm{dia,i}}}. \qquad (26)$$

where u is defined as Φ/F_0 and $\mathbf{v}_P = (\mathbf{e}_\phi \times \mathbf{v}) \times \mathbf{e}_\phi$ denotes the poloidal component of the velocity. In this expression, $\mathbf{E} \times \mathbf{B}$ effects are captured by the first term, and the ion diamagnetic drift velocity by the second one. Given this expression, we can define the approximation space for the $\mathbf{v}_E$ poloidal component of the velocity variable as $\mathcal{V}_{\mathbf{v}_E} = \{\mathbf{v} : \exists u \in H^2$ s.t. $\mathbf{v} =$

$-R\nabla u \times \mathbf{e}_\phi\}$ and according to Ritz–Galerkin method, it is natural to choose the velocity test functions in the poloidal direction in the same space:

$$\mathbf{V}_E^\star = -R\nabla v^\star \times \mathbf{e}_\phi$$

A first version of the reduced MHD model can be obtained using the definition of these spaces. However, for many problems, flows are not purely poloidal and one must take into account flows along the magnetic field lines. Therefore an improved version of the reduced MHD model used in JOREK defines the velocity approximation space by:

$$\mathbf{V} = \mathbf{v}_E + \mathbf{v}_{\mathrm{dia,i}} + \underbrace{v_\parallel \mathbf{B}}_{\equiv \mathbf{v}_\parallel} \qquad (27)$$

This reduced MHD model is thus characterized by a magnetic potential defined by a single scalar function (ψ) and a velocity field defined by the two scalars functions (u and $v_\parallel$). As done for $\mathbf{v}_E$, it will be natural to define the parallel test functions using Ritz–Galerkin recipe as $v^\star\mathbf{B}$ for some scalar $v^\star$.

It is important to point out here that, even if a constant F_0 is used, instead of the function F used in the computation of the Grad–Shafranov equilibrium (i.e. assuming that $RB_\phi = $ const), the reduced model preserves the equilibrium. Indeed, when focussing on the momentum equation, we can prove that all the terms associated with the function F (from the initial Grad–Shafranov equilibrium) are in the kernel of the momentum projectors: $v^\star\mathbf{B}$ and $-R\nabla v^\star \times \mathbf{e}_\phi$. This is a direct consequence of the fact that $R\nabla v^\star \times \mathbf{e}_\phi = R^2\nabla \times (v^\star\nabla\phi)$, $\mathbf{B} \cdot (\mathbf{J} \times \mathbf{B}) = 0$ and $\nabla \times (\nabla F^2/2) = 0$. Note that, although the ansatz (25) neglects toroidal field compression by using a constant F_0 instead of the flux function $F(\psi)$ for RB_ϕ, the Shafranov shift is retained due to the use of the solution of the Grad–Shafranov equation with nonzero FF' as the initial condition for ψ.

Summarizing, the reduced MHD is defined by the magnetic and velocity ansatz given by equations (25) and (26) respectively. The weak form for the reduced MHD equations after integration by parts can be directly derived from the general expressions (17)–(20) taking into account the present definition of the functional spaces. We detail in the sequel the expression of the magnetic potential and momentum equations. For the magnetic potential, it is convenient to use $A^\star\mathbf{B}$ as test function and we obtain

$$\int \frac{\partial\mathbf{A}}{\partial t} \cdot \mathbf{B}A^\star = -\int (\mathbf{E} - F_0\nabla u) \cdot \mathbf{B}\, A^\star$$

that gives the problem: Find $\psi \in H^2$ such that for any $A^\star \in H^2$ we have:

$$\int \frac{1}{R^2}\frac{\partial\psi}{\partial t}A^\star = -\int \left(\frac{1}{R}[u,\psi] + \frac{\eta}{R^2}(j - j_\star) \right.$$
$$\left. - \frac{F_0}{R^2}\frac{\partial u}{\partial\phi} + \frac{\delta_*}{2\rho}\nabla p \cdot \mathbf{B} \right) A^\star \qquad (28)$$

where we have introduced the Poisson bracket $[f,g] = \nabla f \times \nabla g \cdot \mathbf{e}_\phi$. Note, that the poloidal current component has been

neglected in the resistive term. To establish the momentum equation, we first use the expression of the velocity (27) together with our definition of the gyro-viscous tensor (12) to obtain the equation:

$$\rho \frac{\partial \mathbf{v}}{\partial t} = -\rho \mathbf{v} \cdot \nabla \mathbf{v} - \rho \mathbf{v}_{\mathrm{i}}^* \cdot \nabla \mathbf{v}_{\mathrm{E}} - \nabla p$$
$$+ \mathbf{J} \times \mathbf{B} + \nabla \cdot (\underset{f}{\tau} + \underset{\mathrm{neo}}{\tau}) + \mathbf{S}_{\mathbf{V}}, \qquad (29)$$

where $\mathbf{v} = \mathbf{v}_\| + \mathbf{v}_{\mathrm{E}}$. Now, using successively $v^\star \mathbf{B}$ and $\mathbf{V}_{\mathrm{E}}^*$ as test functions allows to obtain two scalar equations for the parallel and poloidal components of the velocity:

$$\int \rho \frac{\partial \mathbf{v}}{\partial t} \cdot \mathbf{B} \, v^\star = \int \left(\frac{\mathbf{v} \cdot \mathbf{v}}{2} \nabla \rho + \rho \mathbf{v} \times \mathbf{w} - \nabla p \right.$$
$$\left. + \nabla \cdot (\underline{\tau}_f + \underline{\tau}_{\mathrm{neo}}) + \mathbf{S}_{\mathbf{V}} \right) \cdot \mathbf{B} v^\star$$
$$- \int \rho(\mathbf{v}_{\mathrm{i}}^* \cdot \nabla \mathbf{v}_{\mathrm{E}}) \cdot \mathbf{B} v^\star + \mathrm{BT}_{v_\|}, \quad (30)$$

$$\int \rho \frac{\partial \mathbf{v}}{\partial t} \cdot \mathbf{V}_{\mathrm{E}}^* = \int \left(\frac{\mathbf{v} \cdot \mathbf{v}}{2R^2} \nabla(R^2 \rho) + \rho \mathbf{v} \times \mathbf{w} \right.$$
$$\left. - \nabla p + \nabla \cdot (\underline{\tau}_f + \underline{\tau}_{\mathrm{neo}}) + \mathbf{S}_{\mathbf{V}} \right) \cdot \mathbf{V}_{\mathrm{E}}^*$$
$$- \int \rho(\mathbf{v}_{\mathrm{i}}^* \cdot \nabla \mathbf{v}_{\mathrm{E}}) \cdot \mathbf{V}_{\mathrm{E}}^* + \int v^* \mathbf{B} \cdot \nabla j + \mathrm{BT}_{v_{\mathrm{E}}},$$
$$(31)$$

where we have used the relation $\mathbf{v} \cdot \nabla \mathbf{v} = \nabla \frac{\mathbf{v} \cdot \mathbf{v}}{2} - \mathbf{v} \times \mathbf{w}$ introducing the vorticity $\mathbf{w} \equiv \nabla \times \mathbf{v}$ and $j \equiv -R J_\phi$. The equations for density and temperature are similar to the full MHD context, but with the prescribed velocity and magnetic field expressions

$$\int \frac{\partial \rho}{\partial t} \rho^\star = -\int \left(\nabla \cdot (\rho \, \mathbf{v}) + \mathbf{v}_{\mathrm{i}}^* \cdot \nabla \rho - S_\rho \right) \rho^\star$$
$$- \int \underline{D} \nabla \rho \cdot \nabla \rho^\star + \mathrm{BT}_\rho, \qquad (32)$$

$$\int \frac{\partial p}{\partial t} p^\star = -\int (\mathbf{v} \cdot \nabla p + \gamma p \nabla \cdot \mathbf{v}$$
$$- (\gamma - 1)(\underline{\tau}_f + \underline{\tau}_{\mathrm{neo}}) : \nabla \mathbf{v} - S_p) \, p^\star$$
$$- \int \underline{\kappa} \nabla T \cdot \nabla p^\star + \mathrm{BT}_p \qquad (33)$$

To derive the final formulation of the pressure equation, gyro-viscous cancelation assumptions have been used. The boundary integrals appearing after the integration by parts are indicated by the symbol BT and defined as

$$\mathrm{BT}_{v_\|} = -\oint v^* \frac{\rho \mathbf{v} \cdot \mathbf{v}}{2} \mathbf{B} \cdot \mathrm{d}\mathbf{S} \qquad (34)$$

$$\mathrm{BT}_{v_{\mathrm{E}}} = -\oint v^* \left(\nabla \left(\frac{R^2 \rho \mathbf{v} \cdot \mathbf{v}}{2} \right) - \mathbf{J} \times \mathbf{B} \right) \times \nabla \phi \cdot \mathrm{d}\mathbf{S}$$
$$(35)$$

$$\mathrm{BT}_\rho = \oint \rho^\star \underline{D} \nabla \rho \cdot \mathrm{d}\mathbf{S} \qquad (36)$$

$$\mathrm{BT}_p = \oint p^\star \underline{\kappa} \nabla T \cdot \mathrm{d}\mathbf{S}. \qquad (37)$$

Note that in this derivation, once the ansatz and projection functions are defined, there are no approximations on geometry. I.e., the reduced MHD derived here is not an aspect ratio expansion of the full MHD model.

These equations involve some high order derivatives whose computations can be alleviated by the introduction of intermediate variables: the toroidal current density (j) and toroidal vorticity (ω) satisfying the following partial differential equations

$$j = \Delta^* \Psi \equiv R^2 \nabla \cdot \left(\frac{1}{R^2} \nabla_{\mathrm{pol}} \Psi \right), \qquad (38)$$

$$\omega = \Delta_{\mathrm{pol}} u \equiv \nabla \cdot \nabla_{\mathrm{pol}} u. \qquad (39)$$

Here, $\nabla_{\mathrm{pol}} u$ denotes the gradient in the R–Z plane. The reduced MHD base model consists of seven scalar physical quantities as variables, see table 1. Five variables are evolved in time ('five field model'), while ω and j are coupled to u and Ψ by definition equations[33]. The evolution and definition equations are solved simultaneously at every time step in a fully implicit numerical scheme (see section 3.2).

2.3.2. Boundary conditions. Boundary conditions can be set in a flexible way. By default, all variables are kept constant in time on the computational domain boundary, wherever the latter is aligned to a flux surface (Dirichlet). Where flux surfaces are intersecting the boundary (e.g., in the divertor region or for grids extended to the true physical wall) sheath boundary conditions are applied as commonly done in divertor physics codes [66]. The poloidal flux, current density, electric potential, and vorticity are kept fixed at the boundary, while the parallel velocity is forced to be equal to the ion sound speed. For the density no Dirichlet condition is forced and the boundary term (36) is not included. Not including the latter boundary term in the finite element method naturally implies that $\underline{D} \nabla \rho \cdot \mathbf{n} = 0$ and therefore the perpendicular ion flux to the boundary is purely convective $\Gamma \cdot \mathbf{n} = \rho \mathbf{V} \cdot \mathbf{n}$. The evolution of the boundary temperatures are constrained by the following BCs for the normal ion and electron heat fluxes to the boundary

$$\mathbf{q}_{\mathrm{i}} \cdot \mathbf{n} \equiv \left(\frac{\rho}{2} \mathbf{V} \cdot \mathbf{V} + \frac{\gamma}{\gamma - 1} \rho T_{\mathrm{i}} \right) \mathbf{V} \cdot \mathbf{n} - \frac{\kappa_{\mathrm{i}}}{\gamma - 1} \nabla T_{\mathrm{i}} \cdot \mathbf{n}$$
$$= \gamma_{\mathrm{sh,i}} \rho T_{\mathrm{i}} \mathbf{V} \cdot \mathbf{n}, \qquad (40)$$

$$\mathbf{q}_{\mathrm{e}} \cdot \mathbf{n} \equiv \frac{\gamma}{\gamma - 1} \rho T_{\mathrm{e}} \mathbf{V} \cdot \mathbf{n} - \frac{\kappa_{\mathrm{e}}}{\gamma - 1} \nabla T_{\mathrm{e}} \cdot \mathbf{n}$$
$$= \gamma_{\mathrm{sh,e}} \rho T_{\mathrm{e}} \mathbf{V} \cdot \mathbf{n}, \qquad (41)$$

where $\gamma_{\mathrm{sh,i}} \sim 2$–3 and $\gamma_{\mathrm{sh,e}} \sim 5$–6 are the ion and electron sheath transmission factors [66]. For the single temperature

---

[33] Via the definition equations, ω and j are projected to the G^1 continuous Bezier basis, while expressing them directly in terms of u and Ψ would correspond to a discontinuous representation.

Table 1. Scalar physics variables of the base reduced MHD model. ω and j are derived variables, connected to u and Ψ by definition equations.

Symbol	Description	
Ψ	Poloidal magnetic flux	$= R\mathbf{A} \cdot \mathbf{e}_\phi$ with $\mathbf{A}$ the vector potential
u	Velocity stream function	$= \Phi/F_0$ with Φ the electric potential
j	Toroidal current density	$= -R\mathbf{j} \cdot \mathbf{e}_\phi = \Delta^*\Psi$
ω	Toroidal vorticity	$= \Delta_{\text{pol}}u$
ρ	Mass density	$= n_e m_{\text{ion}}$ for singly charged ions
T	Temperature	$\equiv T_e + T_i$ in the single temperature model
$v_\parallel$	Parallel velocity	$= \mathbf{v}_\parallel \cdot \mathbf{B}/B^2$

model ($T_e = T_i$), the two latter expressions can be added to find the total heat flux equation

$$\mathbf{q} \cdot \mathbf{n} \equiv \left(\frac{\rho}{2}\mathbf{V} \cdot \mathbf{V} + \frac{\gamma}{\gamma - 1}\rho T \right) \mathbf{V} \cdot \mathbf{n} - \frac{\kappa}{\gamma - 1}\nabla T \cdot \mathbf{n}$$
$$= \gamma_{\text{sh}}\rho T_e \mathbf{V} \cdot \mathbf{n}, \tag{42}$$

where γ_{sh} is the total sheath transmission factor that has typical values of 7–8. The latter BCs are expressed in the form $\underline{\kappa}\nabla T \cdot \mathbf{n} = -(c_b - 1)\rho T\mathbf{V} \cdot \mathbf{n}$ and replaced in the boundary term (37) in order to implement them as natural BCs. For the electrons $c_{b,e} = (\gamma - 1)(\gamma_{\text{sh,e}} - 1)$, for the ions $c_{b,i} = (\gamma - 1)(\gamma_{\text{sh,i}} - \gamma - 1)$ and the total heat flux $c_b = (\gamma - 1)(\gamma_{\text{sh}}/2 - \gamma/2 - 1)$. Note that sheath boundary conditions are applied on the whole boundary of the computational domain if grids extended to the physical first wall are used (see section 3.1.2). In case of free boundary simulations, the Dirichlet condition on the plasma current density and poloidal flux is removed, and a natural boundary condition is implemented instead, like described in section 2.9. Further extensions of the boundary conditions have been developed for particular applications, e.g., a limitation of the current density to the ion saturation current [67].

2.3.3. Properties of the reduced MHD model. Since a significant number of reduced MHD models with very different properties have been proposed in literature, some confusion exists regarding their capabilities. We explain a few key features of our reduced MHD model in the following. A recent discussion of reduced and full MHD models is also provided by reference [68], and section 2.15 shows reduced MHD models for stellarator configurations yet to be implemented, including a detailed discussion of the conservation properties.

In reference [64], it is shown that the model implemented in JOREK is energy conserving as consequence of the full MHD being energy conserving and the ansatz based approach being used. This is strictly valid only for the single-fluid model, where diamagnetic drift effects are excluded, since the gyro-viscous cancelation is not exactly energy conserving [65, 69]. The non-conservation introduced by gyro-viscous cancelation is of the order of $\delta^* p_m/(\rho_0 v_A)$, where p_m and v_A denote the magnetic pressure and Alfvén velocity, respectively. As long as δ^* is small, the error should be acceptable. The energy conservation test shown in the right panel of figure 17 supports a good energy conservation also when diamagnetic drift effects

are included. Note that, even without diamagnetic drift, energy conservation is only exact in the limit of continuous time, and introducing a finite time step also introduces a small error in the energy conservation. Another source of error is the truncation of the toroidal Fourier series (see section 3.1.3). More formally, for a simplified version of the reduced MHD model, it has been shown in reference [53] that reduced MHD models are a valid approximation of the full MHD model, i.e., the solutions of the full MHD system converge to the solutions of an appropriate reduced model.

Linear momentum is not exactly conserved locally by the reduced MHD models even in the continuous limit. This becomes obvious when one acknowledges that each of the three Cartesian components of the momentum equation govern momentum conservation in the respective Cartesian direction, and all three must be satisfied individually in order for linear momentum to be conserved. However, there are only two scalar velocity variables, u and $v_\parallel$, in the reduced MHD model, and only two scalar equations governing these variables, namely (30) and (31), remain in the reduced model. Thus, it is impossible to locally conserve all three components of linear momentum in reduced MHD, except for some special cases. Globally, the z-component of momentum is always conserved, as can be seen by letting $v^* = \ln R$ (which gives $\mathbf{V}_E^* = \nabla z$) in equation (31). One can also ensure global conservation of the x and y components of linear momentum by excluding the $n = 1$ term from the Fourier series, which forces the x and y components of the integrated total momentum to zero due to symmetry. This topic is considered in more detail in reference [70]. The full MHD model is conserving momentum exactly on the equation level, comparisons presented in section 4 between both models provide confidence to some degree that the introduced error does not affect linear and non-linear dynamics significantly.

The presented reduced MHD model satisfies $\nabla \cdot \mathbf{J} = 0$. In fact it can be shown that equation (31) is identical to the weak form of $\nabla \cdot \mathbf{J} = 0$. This is demonstrated by applying the cross product $\times \nabla\phi$ to equation (29) in order to obtain the poloidal current density

$$F_0 \mathbf{J}_{\text{pol}} = -j\mathbf{B}_{\text{pol}} - R^2 \nabla p \times \nabla\phi$$
$$- R^2 \left(\rho\frac{\partial\mathbf{v}}{\partial t} + \rho\mathbf{v} \cdot \nabla\mathbf{v} + \rho\mathbf{v}_{\text{dia, i}} \cdot \nabla\mathbf{v}_E \right.$$
$$\left. - \nabla \cdot (\underline{\tau}_f + \underline{\tau}_{\text{neo}}) - \mathbf{S}_V \right) \times \nabla\phi \tag{43}$$

10

Then using the latter expression in $\int v^* \nabla \cdot (\mathbf{J}_\phi + \mathbf{J}_{\text{pol}}) = 0$ and applying integration by parts, equation (31) is recovered. As it can be inferred from equation (43), even if the toroidal field is fixed in time, the poloidal currents exist in this model and evolve according to the momentum equation and conservation of current. The ansatz (25) together with the projection operator $e_\phi \cdot \nabla \times$ projects out, i.e. removes, the poloidal currents from the system of equations. This does not imply that the poloidal currents are neglected in the model, but rather their contribution to the toroidal field is dropped. The poloidal currents can be calculated, *a posteriori*, from (43). As mentioned above, the reduced ideal MHD momentum equation is consistent with the Grad–Shafranov equation, even in the absence of poloidal currents. The force balance $\mathbf{J} \times \mathbf{B} = \nabla p$ appears as $[\psi, j] - [R^2, p]$ in the momentum equation. Substituting $j = FF'(\psi) + R^2 p'(\psi)$ shows that the two terms in the pressure balance cancel. The benchmarks of vertical displacement event (VDE) simulations between the JOREK reduced MHD model and the M3D-C^1 and NIMROD full MHD models shows the validity of this approach: the agreement regarding plasma dynamics and halo currents is very good (see section 6.3).

As mentioned above, reduced MHD models aim to eliminate fast magnetosonic waves, the fast propagation of which can impose restrictive CFL limits for explicit methods and significantly increase the stiffness of the problem for implicit methods. In the JOREK reduced MHD model presented here, the fast magnetosonic waves are eliminated by the velocity ansatz (27). In reference [71], it is shown that any velocity field can be decomposed into an $\mathbf{E} \times \mathbf{B}$ term, a field-aligned flow term and a perpendicular fluid compression term[34], which are responsible for Alfvén waves, slow magnetosonic and fast magnetosonic waves, respectively. The $\mathbf{E} \times \mathbf{B}$ and field-aligned flow terms are the first and third terms, respectively, in the velocity ansatz (27), whereas the perpendicular fluid compression term is not present in the ansatz.

Finally, it is important to note that the reduced MHD model presented here cannot correctly reproduce pressure-driven modes under all circumstances. In particular, the 1/1 internal kink mode at nonzero β is affected. As shown in reference [68], the term associated with fast magnetosonic waves in the energy functional can be written as

$$\frac{1}{2} \int_V \left| \delta B_\parallel + \delta p / B_0 \right|^2 \, \mathrm{d}V,$$

where $\mathbf{B} = \mathbf{B}_0 + \delta \mathbf{B}$ and $p = p_0 + \delta p$; the quantities with a '0' subscript correspond to the equilibrium, and quantities with a 'δ' prefix are perturbations. Since in our reduced MHD model, we have set $\mathbf{e}_\phi \cdot \delta \mathbf{B} = 0$, we have $\delta B_\parallel = \mathbf{B}_0 \cdot \delta \mathbf{B} / B_0 \approx 0$, whereas in full MHD simulations, one often finds that $B_0 \delta B_\parallel \approx -\delta p$. Thus, in the case of pressure-driven instabilities, the term above contributes a stabilizing effect due to

the δp contribution not being cancelled by $B_0 \delta B_\parallel$ [68]. Traditional ballooning modes in the plasma edge are an exception to this rule. As shown in reference [68], the Mercier criterion is modified as $4\epsilon\alpha(1 - q^2) > s^2 q^2 + \alpha^2$, where $\alpha = -2q^2 R_0 p' / B_0^2$ is the ballooning parameter and $s = rq'/q$ is the shear. Since $q \gg 1$ and $\alpha \sim 1$ in the plasma edge, ballooning modes are largely unaffected, which can be seen in various benchmarks (section 4.5). In figure 19, the linear growth rates for the internal kink mode are shown. For β near zero, the mode is mostly current driven, and the stabilizing effect discussed above is negligible. However, as β increases, the accuracy of reduced MHD quickly deteriorates. Alternate reduced MHD models, such as that by Kruger *et al* [54], can better capture most pressure-driven modes due to the incorporation of the constraint $B_0 \delta B_\parallel = -\delta p$ into their model.

2.3.4. Related models and extensions. The various extensions available for the reduced MHD model are described in the following, sections 2.4–2.11. Note, that also a simplified version of the reduced MHD model is available, where the parallel momentum equation and the variable $v_\parallel$ have been eliminated, while the rest of the description shown above remains unchanged. This corresponds to a drop of slow magneto-sonic waves. The full MHD model of JOREK is described briefly in section 2.12. Formulations of reduced as well as full MHD appropriate for stellarators presently being implemented are shown in section 2.15.

2.4. External magnetic perturbations

For simulations of RMPs used typically for ELM mitigation or suppression, a 3D poloidal flux perturbation at the boundary of the computational domain can be ramped up during the simulation [50]. The perturbation needs to be pre-calculated by an external code (vacuum assumption for the boundary condition). This approach has widely been used for previous studies of RMPs (see section 5.4) with the drawback that the magnetic field perturbation at the computational boundary cannot evolve consistently. Using the free boundary extension (see section 2.9), RMPs can alternatively be described fully consistently from 3D coils [72].

2.5. Separate electron and ion temperatures

An extension for treating electron and ion temperatures separately [73] introduces one additional variable to evolve both temperatures independently in time. In particular, different parallel heat diffusion coefficients can be used for the species, allowing to capture the temperature evolution across an ELM cycle more accurately. The parallel heat conductivity does not only influence the non-linear evolution of the plasma considerably, but also affects linear stability properties (neglected in most stability codes).

2.6. Neutrals

A model extension is available to include a neutral particle fluid in the simulations, a development originally started in reference [74]. The present version was derived and implemented

---

[34] The perpendicular fluid compression term is mostly responsible for plasma compressional motion orthogonal to the background vacuum field, but some such compression is allowed already by the $\mathbf{E} \times \mathbf{B}$ term. See reference [71] and the references therein for more detailed discussion.

in reference [75]. One additional physics variable was introduced to describe the distribution of neutrals across the computational domain. In this model, the neutral transport is purely diffusive. Ionization and recombination terms, as well as radiative loss terms are implemented. Recycling boundary conditions at the divertor targets have recently been implemented [76, 77]. The model is used for deuterium massive gas injection (MGI) or shattered pellet injection (SPI) simulations (see section 6.2) as well as detachment studies (see section 5.6). A kinetic treatment of neutrals is also possible using the framework described in section 2.10 with first applications on the way.

2.7. Impurities

For the modelling of impurities, several options exist. As a particularly simple model to incorporate some impurity effects, a radiative loss term can be switched on in the reduced MHD model with neutrals (section 2.6). Losses are then calculated under the assumption of a spatially and temporally constant background impurity distribution with a prescribed radiative cooling rate.

For a more realistic description, a model exists where impurities are treated as an additional fluid species [78–80]. This model is applied to massive gas injection or SPI simulations (see section 6.2), but also to radiative collapse simulations (see section 6.1.2). One additional variable is introduced to describe the impurity density distribution. All impurities are assumed to be convected together with the main plasma independently of the charge state, and the impurities are assumed to be in coronal equilibrium. The latter assumption may lead, at least in certain cases, to an underestimation of energy dissipation by impurities. For example, for an axisymmetric benchmark case on impurity dynamics [81], JOREK (with its coronal equilibrium model) predicts a roughly two times slower thermal collapse than M3D-C^1 and NIMROD which have a more advanced model tracking the density of each impurity charge state. The coronal equilibrium assumption also results in an instantaneous change in the ionization state according to the electron temperature, resulting in difficulties in treating the ionization energy and the corresponding recombination radiation which would not be present in a self-consistently evolving non-equilibrium model. To avoid such artificial recombination radiation, we currently treat the ionization energy as a potential energy. A more advanced model, going beyond the coronal equilibrium assumption, is presently in preparation (see section 2.15).

Also kinetic particles (section 2.10) can be used to describe impurities. This has already been applied to study the transport of tungsten during ELM crashes (see section 5.1) as well as the sputtering and SOL transport of tungsten (see section 6.1.2).

2.8. Pellets

Several pellet ablation models are available in JOREK both for pellets consisting of the same material as the main plasma ('deuterium pellets') and for pellets consisting of a different

material ('impurity pellets', e.g., argon or neon). These ablation models are combined with the neutrals model (section 2.6) or the impurity model (section 2.7), respectively.

For the particle source corresponding to pellet ablation, scaling laws from various neutral gas shielding type of models in a Maxwellian plasma are implemented [82–85]. The main idea behind these models is that the ablation rate naturally adapts such that the incoming heat flux from the ambient plasma is almost fully absorbed by the ablation cloud surrounding the pellet. Literature provides scaling laws for ablation rates for various pellet materials (including mixtures) which have been obtained by fitting numerical results of gas dynamics simulations [83–85].

Ablated atoms are deposited via a volumetric source term of the form:

$$
\begin{aligned}
S_n \propto &\left[0.5 - 0.5 \, \tanh \left(\frac{(R - R_{\mathrm{p}})^2 + (Z - Z_{\mathrm{p}})^2}{\delta r_{\mathrm{c}}} \right) \right] \\
&\times \left[0.5 - 0.5 \, \tanh \left(\frac{\phi - \phi_{\mathrm{p}}}{\delta \phi_{\mathrm{c}}} \right) \right],
\end{aligned} \tag{44}
$$

where R_{p} and Z_{p} are the pellet location and δr_{c} and ϕ_{c} characterize the poloidal and toroidal extension of the ablation cloud. The parameters δr_{c} and $\delta \phi_{\mathrm{c}}$ determine the width of smoothing of the source profile in poloidal and toroidal direction. The pellet is presently assumed to move along a straight line with constant velocity, and its particle content (and physical size) is evolved according to the ablation. The toroidal extension $\delta \phi_{\mathrm{c}}$ of the ablation cloud in simulations is typically far larger than in reality due to limited toroidal resolution, but tests shown in reference [86] for a deuterium pellet found that for a sufficiently small $\delta \phi_{\mathrm{c}}$, JOREK results converge, i.e. MHD dynamics becomes independent of $\delta \phi_{\mathrm{c}}$. For impurity pellets, the same may however not be true because of the radiative loss term, which scales like $n_{\mathrm{imp}} n_{\mathrm{e}}$, i.e. like n_{imp}^2 in regions where the impurity density n_{imp} is large, such that the total power radiated in the ablation cloud scales like $1/\delta \phi_{\mathrm{c}}$. For SPI simulations [87], the model described above is applied for each individual shard. Input parameters allow specifying the shard size distribution, the averaged velocity and velocity spread of the shards.

2.9. Free boundary and resistive walls

Via a coupling [88] to the STARWALL code [89], JOREK is capable of free boundary simulations. In the Greens functions approach applied here, STARWALL discretizes the conducting structures by triangles (thin-wall approximation), while the vacuum region surrounding the plasma and conducting structures is not discretized. The JOREK-STARWALL coupling is then performed via a natural boundary condition at the edge of the JOREK computational domain that replaces the Dirichlet boundary conditions for the poloidal flux and current density. In a boundary integral, that arises from partial integration (see section 3.1.4) of the current definition equation, and that vanishes for fixed boundary simulations, the tangential magnetic field is expressed in terms of the poloidal flux values and the currents in the conducting structures as shown in detail in reference [88]. The BC is a Neumann type condition for the

magnetic vector potential, which results from the analytical solution of the vacuum field given by the Green's functions. In terms of the response matrices, the magnetic field has the form

$$\mathbf{B} \times \mathbf{n} = \underline{\mathbf{M}}_{\mathrm{vac}} \, \mathbf{B} \cdot \mathbf{n} + \underline{\mathbf{M}}_I I,$$

where $\mathbf{n}$ is the normal vector to the boundary, $\underline{\mathbf{M}}_{\mathrm{vac}}$ is the vacuum response matrix and $\underline{\mathbf{M}}_I$ is a matrix that calculates the contribution of the wall and coil currents (I). The evolution of the wall currents is calculated with resistor–inductor circuit equations that arise for each of the discretized wall elements. The 'response matrices', which allow to calculate the evolution of the wall currents and the tangential magnetic field at the JOREK boundary, are calculated by STARWALL only once in the beginning of a simulation. Since they are only dependent on the JOREK grid and wall geometry, response matrices may even be re-used for further simulations if the geometry remains the same. The response matrices are written out by STARWALL into a file and read by JOREK using MPI I/O in both codes. In JOREK, the matrices are used to evolve wall currents in time and to implement the natural boundary condition. The coupling between plasma and wall currents is implemented in a fully implicit way that is entirely consistent with the time evolution of the intrinsic JOREK equations. The dimensionality of the sparse matrix system is not increased in spite of this fully implicit approach compared to fixed boundary simulations since the implicit values of the wall currents are analytically eliminated from the system[35].

JOREK–STARWALL allows to choose a fixed or free boundary mode independently for each toroidal harmonic. This is sometimes used to keep fixed boundary conditions for the axisymmetric $n = 0$ component, while a free boundary treatment is applied to non-axisymmetric $n \neq 0$ components. For simulations, where also the $n = 0$ component is treated to be free, the equilibrium solver has been extended to free boundary cases [88] and has recently been updated for Newton iterations to enhance convergence (using the methods described in reference [90]). Magnetic field coils have been implemented self-consistently in STARWALL, including time varying coil currents and their interaction with conducting structures [91] and recently also arbitrary 3D coils have been implemented in a self-consistent way [72] (see also section 4.8). This allows to include active coils (poloidal field coils, RMP coils, etc) and passive coils (Mirnov coils, saddle coils, etc) consistently in JOREK–STARWALL simulations. A functionality is available also, which allows to create a free boundary equilibrium for a given fixed boundary case, by automatically determining appropriate coil currents [92].

Via the derivation shown in references [93–95], plasma currents flowing directly into conducting structures or out of them (current sharing between plasma and wall), can be treated

consistently with the STARWALL formalism. The respective derivation for JOREK–STARWALL including the interaction of eddy and halo currents are shown in reference [92]. However, the implicit coupling of wall source/sink (halo currents) with the plasma electric potential has not been implemented yet. For walls with a low poloidal path resistance, the usual JOREK BC for the electric potential ($\Phi = 0$) gives the correct distribution of halo currents as demonstrated in [96]. The formalism derived in [93–95] has been implemented as a post-processing tool to calculate wall forces and to visualize the source/sink currents (see figure 2). This tool has been validated as well for 3D walls with holes. A formalism for treating ferromagnetic components in a thin wall model is shown in reference [97] but integration in JOREK–STARWALL hasn't been approached so far.

Some numerical limitations had originally restricted JOREK–STARWALL to moderate toroidal and poloidal resolutions. In particular, STARWALL was originally purely OpenMP parallelized, the coupling terms inside JOREK were treated OpenMP parallel only, and the fairly large and sparse 'response matrices' were duplicated across all MPI ranks in JOREK. Within the project described in reference [99], an MPI parallelization of STARWALL, a hybrid MPI + OpenMP parallelization of the coupling terms in JOREK, parallel input/output for the response matrices in both codes as well as a distributed storage of the response matrices across the MPI ranks were implemented including the distributed matrix–matrix operations. With these developments, high resolution cases are possible now with an excellent performance. For a verification of free boundary simulations, see section 4.8.

2.10. Kinetic particles

For a number of applications, such as the transport and interaction of fast particles, impurities and neutrals with the MHD fluid, the main fluid model(s) in JOREK have been extended with a kinetic particle module. Particles are followed in the time dependent 3D magnetic and electric fields given on the cubic finite element grid. For the fast ions, impurities and neutrals, the well-known Boris scheme is used. The full orbit of the particles is followed both in real (R, Z, ϕ) space and in the local coordinates (s, t, ϕ) of each element. At every particle step, the Boris scheme is applied in (R, Z, ϕ) coordinates including a correction for the toroidal geometry [100]. The updated local element coordinates are found by Newton iterations. Following the particles in real space has the advantage that the crossing of a finite element boundary does not influence the Boris scheme. The change of element of a particle position is handled in the update of the local element coordinates. For the rare case a particle crosses many element boundaries, an efficient RTree algorithm is used to find the particles element. The Boris scheme combined with a higher order interpolation of the magnetic and electric fields in time assures accurate particle trajectories with good energy conservation (see section 4.9). Multiple particle species can be traced within one simulation allowing for example simulations combining neutral deuterium, heavy impurities and fast ions. For the simulation of relativistic REs, a relativistic

---

[35] Note, that the natural boundary condition, however, leads to a less sparse matrix structure for boundary degrees of freedom. Local interactions between neighbouring grid nodes are replaced by global interactions on the boundary. To ensure efficient load balancing also for such simulations, the domain decomposition is slightly adapted compared to fixed boundary simulations.
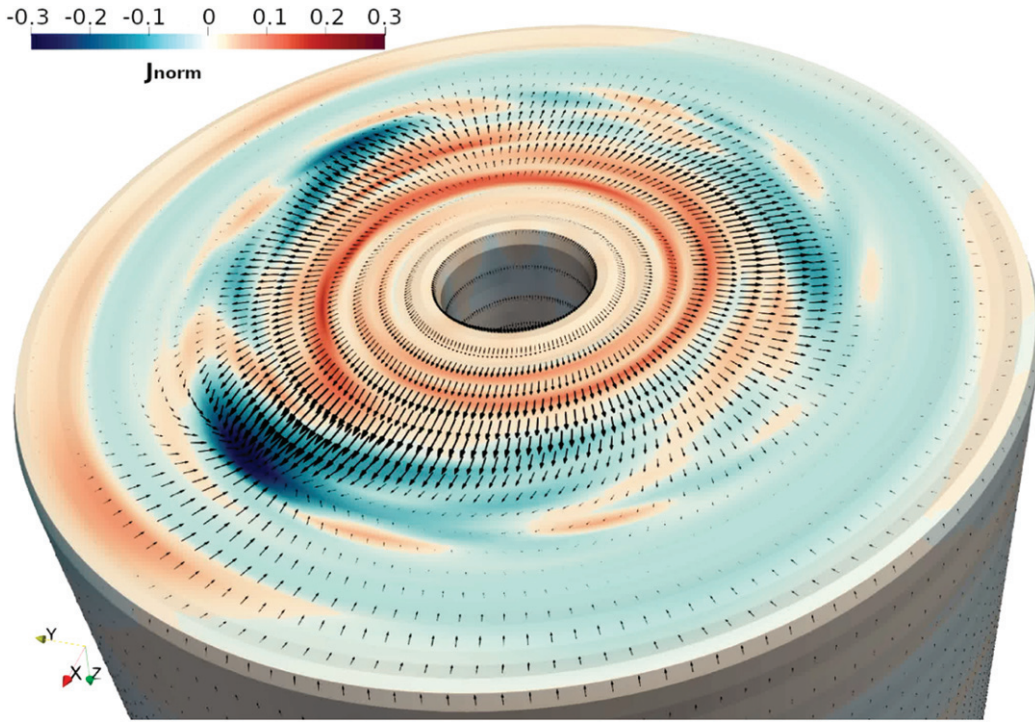
Figure 2. Bottom view of the thin wall used for the 3D VDE benchmark case published in reference [98] showing the distribution of source/sink (halo) currents (black arrows). The colours indicate the perpendicular current density into the wall respectively out of it (in a.u.). The BC for the electric potential was $\Phi = 0$ here such that the wall acts ideally conducting for the halo currents.

particle tracer (both full orbit or guiding centre) was implemented [101, 102] with excellent energy and momentum conservation properties (see section 4.9). This model was applied to study various aspects of RE dynamics with a test particle approach (section 6.4). In addition to the particle following, the charge state of the particles and the ionisation, recombination and radiation rates are consistently evolved in time according to the background fluid properties, using the OPEN-ADAS rate coefficients including all charge states. A model for particle collisions with the background fluid or with projected moments of the kinetic particles, following references [103–105] has been implemented and successfully benchmarked with the cases in references [104, 105]. The collision model includes the thermal force, relevant for the movement of impurities upwards relative to the temperature gradient in on open field lines. To model the main source of impurities, the sputtering of divertor/first wall material by incoming particles and MHD fluid has been implemented using the sputtering yields from references [106–108] and benchmarked with the results from reference [109]. The first application, with test particles, studied the transport of tungsten impurities during (simulations of) ELM crashes (see section 5.1), tungsten sputtering during ELMs and scrape-off layer (SOL) transport of tungsten (see section 6.1.2). For the modelling of the interaction of the MHD fluid and particles, coupling terms have been implemented. These terms appear mostly as additional explicit source terms on the right-hand side of the fluid equations. To describe the excitation of MHD instabilities driven by fast particles, both the coupling schemes based on the pressure and the current have been

implemented. To include the neutral and impurity physics, density, momentum and energy sources resulting from ionisation, recombination, charge exchange and radiation have been added. From the particle distribution, the source terms are calculated by projecting the moments of the particle distribution onto the finite element representation. For the mass density source S_ρ:

$$S_\rho(\mathbf{x}) = \int m_{\mathrm{i}} f(\mathbf{x}, \mathbf{v}) \, \mathrm{d}\mathbf{v} \tag{45}$$

$$\tilde{S}_\rho(\mathbf{x}) = \sum_{ijk} p_{\rho,ijk} H_{ij}(s,t) H_{\phi,k}(\phi) \tag{46}$$

where $p_{\rho,ijk}$ are the finite element expansion coefficients to be determined by the projection. The projection uses the weak form with the same testfunction as the finite element basis functions $v^* = H_{ij} H_{\phi,k}$:

$$\int v^* \tilde{S}_\rho(\mathbf{x}) \, \mathrm{d}\mathbf{x} = \int v^* S_\rho(\mathbf{x}) \, \mathrm{d}\mathbf{x} = \sum_i v^*(\mathbf{x}) \, m w_i \delta(\mathbf{x} - \mathbf{x_i}). \tag{47}$$

Here, the integral over the particles becomes a sum over all particles weighted by the finite element basis functions. The system of equation (47) is factorised only once and solved at particle projection. The projection results can be smoothed by solving instead:

$$\int v^* \left(1 + \alpha \nabla^4\right) \tilde{S}_\rho(\mathbf{x}) \, \mathrm{d}\mathbf{x} = \int v^* S_\rho(\mathbf{x}) \, \mathrm{d}\mathbf{x} \tag{48}$$

Typical values for α are, depending on the application, of the order of 10^{-10} to 10^{-12}. The projection is required at every

time step of the main fluid part. The number of particle steps for each fluid step varies from 100–1000 for fast particles to order 1–10 for slow heavy impurities. The sources are either a time integrated source (typically required for good conservation properties for particle/energy sources), or at one given time (for fast particle pressures).

2.11. Relativistic electron fluid

A relativistic electron fluid model is available [110], which allows to simulate generation, transport, and losses of REs. One additional variable is introduced to describe the spatial distribution of the RE density. Most of the relevant primary and secondary generation mechanisms have been implemented, and successfully been benchmarked against lower-dimensional codes. This fluid approach does not capture some kinetic aspects (i.e., accurate treatment of the energy spectrum), but allows to study the mutual non-linear interaction between REs and MHD. In that sense, the approach is complementary to the RE test particle model described in the previous section, which captures kinetic effects but does not account yet for the back-reaction of the REs to the plasma. The model has so far been applied to the interaction of REs with internal kinks, VDEs, and TMs as well as RE beam termination studies (section 6.4). An extension has very recently been developed, which takes into account the effect of REs onto the radial force balance of the plasma [111]. The implementation of the interaction terms with the impurity fluid is under development.

2.12. Full MHD model

A full MHD model suitable for production simulations is available in JOREK [65, 112, 113] and can be used for many applications. The model was, nevertheless, only used for few physics applications so far, since its final numerically robust implementation has only been completed recently. Several important extensions available for the reduced MHD model have now been implemented in the full MHD model already.

As discussed in section 4, benchmarks of the reduced and full MHD models of JOREK show that the reduced MHD model is capturing the key physics very well under many conditions while reducing computational costs. However, it is also shown that for certain types of instabilities such as the internal kink, it is necessary to use the full MHD model as also has been found in analytical calculations [68]. Furthermore, VDE benchmarks revealed an overall excellent agreement between the reduced MHD JOREK model and full MHD codes (section 4.8). However, the toroidal variation in the plasma current is not reflected.

The full MHD model has recently been extended for sheath boundary conditions and numerical stabilization terms were implemented [114]. The full-MHD physics model implemented in reference [65] includes plasma flows like the reduced-MHD model, with diamagnetic terms, neoclassical poloidal friction, and toroidal rotation. The bootstrap current source has also been included. A neutrals fluid model like used for MGI and SPI disruption simulations, has not been implemented yet and will be addressed in future studies.

2.13. Electrostatic fluid turbulence model

The JOREK code is very flexible for the implementation of many other physical models not only related to MHD. The realistic geometry, global equilibrium obtained from the Grad–Shafranov solver, flux-aligned grid, numerical methods, sparse matrix solver are common for many applications of the JOREK code. An electrostatic turbulence model has been implemented into JOREK which allows to study ITG driven turbulence in realistic tokamak geometry including SOL and X-point. Both fluid [115–117] and full kinetic orbits approaches are under development. Benchmarks with the standard cyclone case [118, 119] for fluid and kinetic approaches have already confirmed that the implementation is capturing the growth rates of the instabilities in simplified configurations accurately [116, 117]. The model has recently been extended to model ITGs in X-point geometry with SOL for realistic JET and COMPASS discharges (see section 7.1).

2.14. Fully kinetic electrostatic model

The particle framework described in section 2.10 has been used to implement a fully kinetic (electrostatic) model (i.e., no fluid part) of the plasma with full orbit ions and adiabatic electrons. In this case, there is only one equation to solve for the electric potential Φ in terms of ion density n_i:

$$N > 0 : \Phi_N(\mathbf{x}) = n_{i,N}(\mathbf{x}) \tag{49}$$

$$N = 0 : \nabla \cdot \left(\frac{m_e T_e}{e^2 B^2} \nabla_\perp \Phi \right) + \left(\Phi - \langle \Phi \rangle \right) = \frac{n_i - n_{i0}}{n_{e0}} \tag{50}$$

where N is the toroidal harmonic, n_{i0} and n_{e0} the initial ion and electron density. Due to the full orbit ions, there is no ion polarisation density in equation (50). This equation is solved as a slightly modified form of the projection operator from equation (47). The full orbit, full-*f* model has been successfully benchmarked against the linear ITG growth rates and frequencies from [119] and against the zonal flow frequencies and damping from [120], see section 7.1. The model can be used with any of the JOREK finite element grids, allowing ITG simulations in arbitrary X-point geometry, including the open field line region (see section 7.1).

2.15. Outlook

Various extensions to the models are in preparation, and only some of them can be mentioned here. Present plans involve a more modular structure of the physics models to simplify combining arbitrary extensions when needed. Sheath boundary conditions will be further refined, in particular allowing for outflows with Mach numbers larger than one. A fully consistent neoclassical model for the plasma resistivity is being implemented.

Two types of impurity models going beyond the coronal equilibrium assumption are under development: a fluid model involving one additional continuity equation per impurity charge state (or bundles of charge states), and a model in

which impurities are treated as a set of particles, the charge of each particle being evolved individually.

Regarding the RE fluid model, it is planned to evolve in time an average parallel momentum of the particles, allowing to capture the effect of the REs onto the radial force balance, which leads to a major-radial shift between the flux surfaces and the RE drift-orbit surfaces in the absence of thermal pressure.

An effort is presently taken to implement an energy conserving current-coupling scheme for the kinetic particles and the MHD fluid allowing to study the energy exchange of MHD modes with super-sonic particles. Related to that, kinetic MHD or simplified gyrokinetic models like described in references [121, 122] are presently evaluated and might be implemented in JOREK.

In preparation of a 3D extension aiming to simulate stellarator plasmas, a hierarchy of reduced and full MHD models with good energy and momentum conservation properties has been derived in a form suitable for stellarators [70, 71]. The reduced model with some further improvements is presently being implemented into JOREK. Further extensions of the numerical methods, in particular regarding spatial discretization and solver, are necessary to achieve well resolved stellarator simulations and are evolving in parallel (see section 3.8).

3. Numerical methods

The JOREK code is largely written in Fortran with a few core routines in C and C++. It has only few library dependencies allowing to port the code easily to new machines. Some aspects of the code development are described in section 3.7. The numerical methods used in the code are shown in the following, in particular the spatial and temporal discretization, numerical stabilization schemes, the construction of the sparse matrix system, the iterative solver with its preconditioner, and the hybrid parallelization of the code are described.

3.1. Spatial discretization

JOREK solves the equations of the respective physics models in weak form on a G^1 continuous 2D isoparametric Bezier finite element grid combined with a toroidal Fourier expansion [2, 40]. The Bezier basis and the finite element grids constructed from them are shown in subsections 3.1.1 and 3.1.2, respectively. The toroidal Fourier expansion is described in subsection 3.1.3.

3.1.1. Bezier basis. Two-dimensional third order Bernstein polynomials are defined by

$$B_{i,j}^3(s,t) = B_i^3(s)\, B_j^3(t) \quad i,j = 0\dots 3, \qquad (51)$$

with

$$B_i^3(s) = \frac{3!}{i!(3-i)!}\, s^i (1-s)^{3-i} \qquad (52)$$

and where $0 \leqslant s \leqslant 1$ and $0 \leqslant t \leqslant 1$ denote the element–local coordinates, which take values of 0 respectively 1 at the four element vertices (see figure 3 and appendix A). All physical
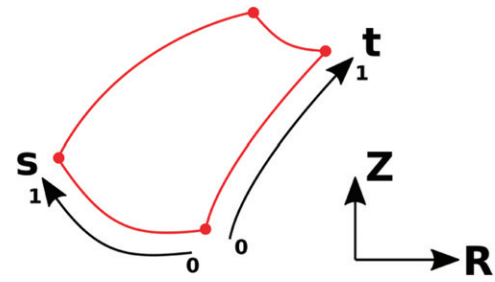
Figure 3. Single 2D Bezier element in *R*–*Z* space. The element–local *s*–*t* coordinate system is also shown. The mapping from *s*–*t* into the *R*–*Z* space is expressed in Bernstein polynomials. The same basis functions are also used for physical variables like the temperature (iso-parameteric).
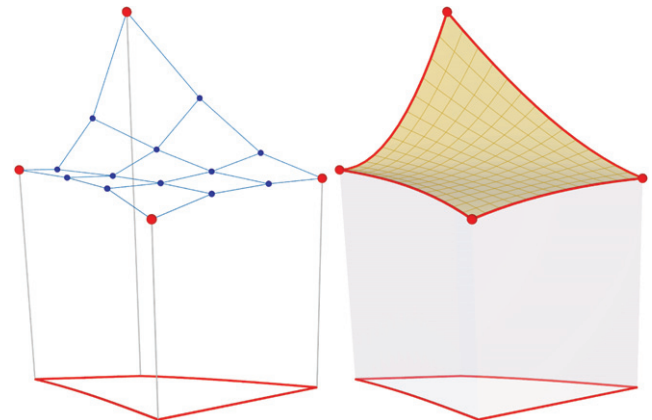
Figure 4. A Bezier element is shown in 3D space. The left plot represent the control points of the elements (blue). The red points are the grid nodes. The right plot depicts the resulting surface. The red structure at the bottom of the plots represents the element in (*R*, *Z*) space. Within each element, a local (*s*, *t*) coordinate system is defined (figure 3). The elevation represents the spatial distribution of a physical variable, e.g., the temperature.

quantities, and also the *R* and *Z* positions of the elements themselves (i.e., the mapping from element local to global coordinates) are expressed in this basis:

$$\mathbf{X}(s,t) = \sum_{i=0}^3 \sum_{j=0}^3 \mathbf{P}_{i,j}\, B_{i,j}^3(s,t) \qquad (53)$$

with $\mathbf{X}$ being an *N*-dimensional vector containing *R*, *Z*, and all interpolated variables of the model (e.g., *R*, *Z*, Ψ, *T*, ρ, etc). Here, $\mathbf{P}_{i,j}$ denotes the Bezier control points in the *N*-dimensional space (see figure 4). For all physical variables, this 2D expression is multiplied with the toroidal Fourier basis (next section), while for *R* and *Z*, axisymmetry is assumed. The coordinates *s* and *t* are orthogonal to the toroidal coordinate ϕ.[36]

---

[36] An ongoing 3D extension of the code will generalize the mapping to be non-axisymmetric, i.e., *R* and *Z* will also be expanded in the Fourier basis. In addition, the orthogonality of *s* and *t* with respect to ϕ will be given up depending on the choice of coordinates. See sections 2.15 and 3.8.

Since first order continuity in real space (G^1), i.e. continuity of the values and the real space gradients, is applied as constraint between neighbouring elements, corresponding control points of adjacent elements need to lie on a common line in the N-dimensional space through the vertex location, which reduces the number of degrees of freedom (see figure 5). Effectively, only four degrees of freedom p_k, u_k, v_k, and w_k remain per node k and for each component of the vector X, which are shared by all elements connected to the respective node. These degrees of freedom are linked to the value, s-derivative, t-derivative, and s–t cross-derivative at the location of the grid node. The relation of the control points with the coefficients p_k, u_k, v_k, and w_k is explicitly shown in reference [2]. Since metric tensor and Jacobian differ between elements, each element e has specific scale factors $d^e_{k,l}$ for each degree of freedom (dof) l at each node k to guarantee G^1 continuity. These scale factors are a geometric grid property and therefore time independent and identical for each physical quantity. For any given quadrangular element e of the mesh, a physical variable, e.g., temperature T, is expanded in the G^1 continuous basis in the following way (the Fourier expansion in toroidal direction is omitted here for clarity):

$$T(s,t)|_e = \sum_{k=1}^{n_{\mathrm{vert}}} \sum_{l=1}^{n_{\mathrm{dofs}}} T_{k,l} H_{k,l}(s,t) \, d^e_{k,l} \tag{54}$$

where $n_{\mathrm{vert}} = 4$ is the number of vertices per element, $n_{\mathrm{dofs}} = 4$ is the number of degrees of freedom per vertex. $H_{k,l}$ are the 16 basis functions written as a product of 1D basis functions in s and t: $H_{k,l}(s,t) = H_k(s)H_l(t)$. The four 1D basis functions $H_k(s)$ are constructed as linear combinations of Bernstein polynomials to satisfy:

- $H_1(s=0)=1$, $H_1'(s=0)=0$, $H_1(s=1)=0$, $H_1'(s=1)=0$

- $H_2(s=0)=0$, $H_2'(s=0)=1$, $H_2(s=1)=0$, $H_2'(s=1)=0$

- $H_3(s=0)=0$, $H_3'(s=0)=0$, $H_3(s=1)=1$, $H_3'(s=1)=0$

- $H_4(s=0)=0$, $H_4'(s=0)=0$, $H_4(s=1)=0$, $H_4'(s=1)=1$

i.e., the well-known cubic Hermite finite elements. This way, the coefficients $T_{k,l}$ become node properties (nodal formulation) which all elements containing the respective node are sharing, while the coefficients $d^e_{k,l}$ corresponding to the same node are different for each element to guarantee the G^1 continuity.

G_1 continuity is not strictly enforced at the grid axis and in the direct vicinity of the X-point(s) in a flux surface aligned grid due to the specific topology at these points (more than four elements share a common node). Although this is typically not an issue in the simulations, an implementation of strict G^1 continuity also at these special points is on its way by locally combining basis functions in an appropriate way (see section 3.8). In a few cases, the higher order FE formulation can produce 'overshoots' leading to zero or negative values of density and temperature, in particular in the presence of strong convection. For that reason, numerical stabilization has
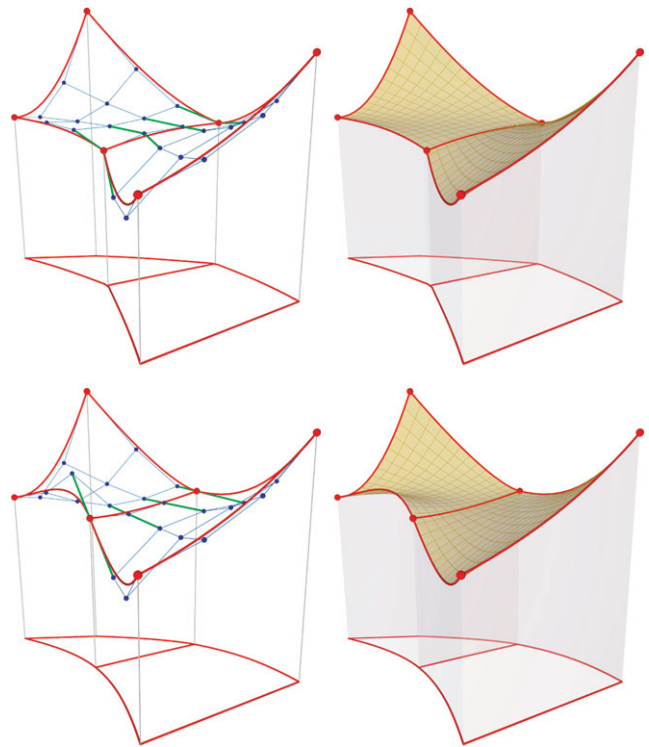
Figure 5. Two neighbouring Bezier elements are shown. The left plots represent the control points of the elements and the right plots depict the resulting surface. In the **plots at the top**, the neighbouring Bezier elements share the control points at the common boundary such that values are continuous across the boundary between the elements (G^0 continuity). However, since the control points around the boundary are not aligned to each other, i.e., green lines exhibit bends, derivatives are not continuous across the element boundary. The **plots at the bottom** correspond to the discretization used in JOREK, where both values and derivatives are continuous across the element boundaries (G^1). Here, the control points at the element boundaries are aligned reducing the degrees of freedom, i.e., the green lines do not exhibit bends.

been implemented as shown in section 3.3. The cubic Bezier element formulation is a generalisation of the isoparametric cubic Hermite elements [39], allowing a local refinement of each element in 2 or 4 sub elements [2].

3.1.2. Finite element grids. All data structures and routines in JOREK are implemented in a general way suitable for unstructured grids. Nevertheless, most grids in use are effectively flux surface aligned grids which are structured within the respective topological domains. Usually, a JOREK simulation is started by calculating the equilibrium on an initial grid, which is not flux surface aligned. Based on this equilibrium, a flux surface aligned grid is then constructed. For numerical accuracy, using such a flux surface aligned grid is in general beneficial (see section 4.2), while the code can be used also with non-aligned grids where appropriate. The grid types available include simple rectangular grids, simple polar grids, polar grids with a rectangular central region. These grids are mostly used as initial grid and only occasionally for the time evolution, e.g., in case of VDEs, where the grid would not remain aligned during the simulation anyways.

Flux surface aligned grid generators are available for plasmas without X-point, with an upper X-point, a lower X-point,
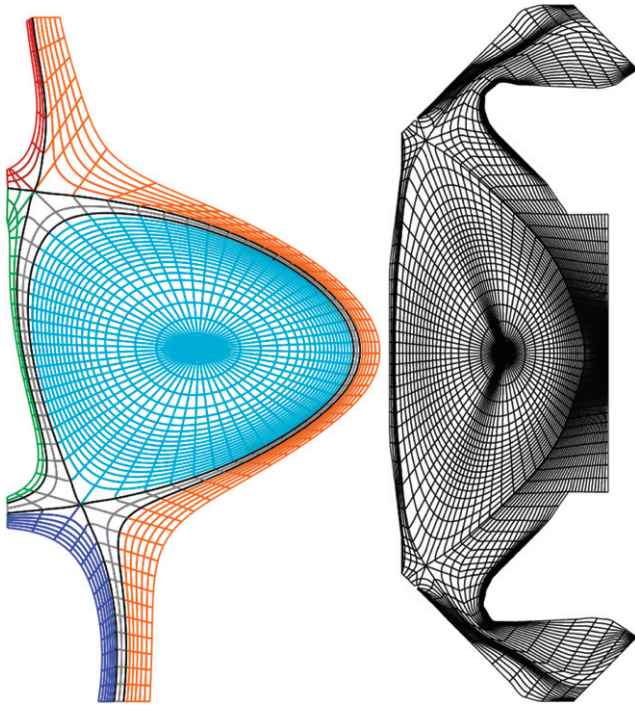
Figure 6. JOREK example grids shown at reduced resolution for clarity. (Left) Flux-surface aligned grid [123] for a TCV equilibrium with two X-points (the lower X-point is active). The part in the closed flux region is shown in light blue, the two separatrices in black, the grid between the separatrices in grey, the lower private flux region in dark blue, the upper private flux region in red, the outer open flux region in orange, and the inner open flux region in green. (Right) Flux surface aligned X-point grid extended to the true first wall (using the methods described in reference [124]) for an MAST Upgrade case with Super-X divertor.

or two X-points. Additionally, an extension of grids to the true first walls is possible as described in [124]. Grid patches (which are not flux-surface aligned any more) are added in the far SOL to cover the space between a flux-surface aligned X-point grid and plasma facing components. Figure 6 contains two examples for a double X-point grid and a double X-point grid extended to the true physical wall (both are shown at reduced resolutions for clarity).

3.1.3. Toroidal Fourier series. As mentioned before, the poloidal discretization in 2D Bezier finite elements is combined with a toroidal Fourier expansion for the physical variables. The toroidal real Fourier harmonics included in a JOREK simulation are selected by the number of harmonics n_{tor}, where cosine and sine components are counted separately—since $n = 0$ does not have a sine component, n_{tor} must be an uneven number. In addition, the periodicity n_{period} can be specified. The latter parameter allows to enforce periodicity in the torus after a fraction $1/n_{period}$ of a full toroidal turn. This effectively means, that instead of $n = 0, 1, 2, \ldots$, the harmonics $n = 0, 1 \times n_{period}, 2 \times n_{period}, \ldots$ are included in a simulation. The number of toroidal planes n_{plane} for the integration along ϕ direction needs to be chosen large enough to avoid aliasing. Typically, $n_{plane} \geqslant 2(n_{tor} - 1)$ is sufficient, i.e., 4 planes per period. While fast Fourier transforms (FFTs) performed during the matrix construction using the FFTW library

(see section 3.4) can handle arbitrary integer numbers, best performance is usually achieved, when n_{plane} is taken to be a power of two.

3.1.4. Weak form. As a consequence of using the finite element approach, the model equations are solved in weak form such that each equation is multiplied by a test function from the same space as the basis functions, and then integrated over the volume V of the computational domain (Galerkin method). Since the resulting equation needs to be fulfiled for each test function, many linear relations are obtained, which form the rows of the matrix system (see section 3.4 for more details on the matrix construction). In case of operators involving higher order derivatives, partial integration is performed e.g.,

$$\int_V \mathrm{d}V a \nabla \cdot \mathbf{b} = -\int_V \mathrm{d}V \nabla a \cdot \mathbf{b} - \oint_A \mathrm{d}A a \mathbf{b} \cdot \mathbf{n}. \tag{55}$$

Here, a and $\mathbf{b}$ are arbitrary scalar respectively vector expressions, A denotes the boundary of the computational domain and $\mathbf{n}$ the normal vector to that boundary. Depending on boundary conditions, some of the boundary integrals vanish, otherwise they are implemented. In particular cases, relations known from physical properties of the system may be inserted into the boundary integrals. In such a case, the boundary integral is expressed in terms of other variables and natural boundary conditions are obtained[37]. Such a natural boundary condition is used, for instance, for the JOREK–STARWALL coupling, where the STARWALL expression for the magnetic field tangential to the boundary is plugged into the boundary integral of the current definition equation (see section 2.9).

3.2. Temporal discretization

The physical equations expressed in the form $\partial \mathbf{A}(\mathbf{u})/\partial t = \mathbf{B}(\mathbf{u}, t)$, where $\mathbf{u}$ denotes the vector of physical variables, are discretized in time as

$$\left[(1 + \xi)\left(\frac{\partial \mathbf{A}}{\partial \mathbf{u}}\right)^n - \Delta t \theta \left(\frac{\partial \mathbf{B}}{\partial \mathbf{u}}\right)^n \right] \delta \mathbf{u}^n$$
$$= \Delta t \mathbf{B}^n + \xi \left(\frac{\partial \mathbf{A}}{\partial \mathbf{u}}\right)^n \delta \mathbf{u}^{n-1}, \tag{56}$$

where the second order linearised Crank–Nicolson scheme is selected by $(\theta, \xi) = (1/2, 0)$ and the second order BDF2 Gears scheme is selected by $(\theta, \xi) = (1, 1/2)$. The first order implicit Euler method corresponds to $(\theta, \xi) = (1, 0)$ but is normally not used in practice. Here, $\delta \mathbf{u}^n = \mathbf{u}^{n+1} - \mathbf{u}^n$ denotes the change of the variables from time step n to $n + 1$. The definition equations for current and vorticity are directly implemented to be satisfied at time point $n + 1$, the above scheme does not need to be applied.

Details, and the time stepping scheme suitable for a more general form of equations is shown in appendix C. The linearization involved in the above time stepping can be replaced

---

[37] Since natural boundary conditions are not enforced, it needs to be checked whether a simulation actually fulfils the original relation. If they are not satisfied well enough, smaller time steps are typically required.

by Newton iterations with a beneficial effect onto non-linear stability in some cases, as demonstrated in reference [64], however this is not implemented in the present code version as these earlier tests showed that larger time steps are possible, however at an increased computational cost and memory consumption such that the overall benefits were limited.

3.3. Numerical stabilization

When the system considered is dominated by convection, the finite element method may suffer from numerical instabilities as illustrated on the left-hand side of figure 7. The variational multi-scale (VMS) is a general framework to avoid these numerical artefacts [125, 126]. The main purpose of the VMS is to take into account the effect of the unresolved scales onto the numerically resolved scales. The consequence is to complement the physical model by a proper numerical dissipative term that scales with the mesh size. The SUPG and the Taylor–Galerkin [127] methods are subsets of the VMS. Let us explain with the mass conservation equation how a simplified VMS contribution is added to the weak form. Assume the initial equation is

$$\frac{\partial \rho}{\partial t} + \mathbf{V} \cdot \nabla \rho = -\rho \nabla \cdot \mathbf{V} + S_\rho$$

In order to stabilize the discretization of the convection term $\mathbf{V} \cdot \nabla \rho$, a dissipation term is added to the weak form to remove numerical artefacts without altering the physical dynamics (figure 7). With this numerical dissipation, the equation that is numerically solved now turns into the following form

$$\frac{\partial \rho}{\partial t} + \mathbf{V} \cdot \nabla \rho = -\rho \nabla \cdot \mathbf{V} + \nabla \cdot (\underline{D}_{\mathrm{vms}} \nabla \rho) + S_\rho.$$

Here, the stabilization tensor $\underline{D}_{\mathrm{vms}}$, in the context of magnetized plasma, is designed such as to take into account the strong anisotropy of the flow

$$\underline{D}_{\mathrm{vms}} = D_{\parallel,\mathrm{vms}}(\mathbf{b} \otimes \mathbf{b}) + D_{\perp,\mathrm{vms}}(\underline{I} - \mathbf{b} \otimes \mathbf{b})$$

where

$$D_{\parallel,\mathrm{vms}} = \tau_\parallel^* (\mathbf{V} \cdot \mathbf{b})^2 \quad \text{and}$$

$$D_{\perp,\mathrm{vms}} = \tau_\perp^* (\mathbf{V} \cdot \mathbf{V} - (\mathbf{V} \cdot \mathbf{b})^2).$$

The scaling factors $\tau_\parallel^*$ and $\tau_\perp^*$ are positive and of the order of the mesh size. More precisely, scaling factors correspond to the ratio of the mesh size over the maximum speed of the acoustic waves and thus have units of time. This scaling ensures that, under mesh refinement (acoustic speeds do not change as the same physics is considered), the VMS numerical dissipation asymptotically vanishes. Therefore, the modified equation numerically solved is consistent with the initial physical model. For the Taylor–Galerkin approach [127, 128] $\tau_\parallel^* = \tau_\perp^* = \delta t/2$. Sometimes, we directly fix $D_\parallel$ and $D_\perp$ to values that scale with the mesh size and the estimated background turbulence. Similar stabilization is applied to the other equations. For the full MHD model (section 2.12), a more general stabilization has been also developed [129], where the
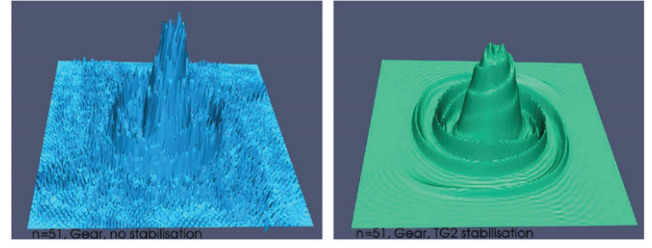
Figure 7. Example of the effect of the Taylor–Galerkin stabilisation (TG2) for a simplified model for the vorticity evolution (reference [130]) in 2D on a grid of 50 by 50 cubic finite elements.

hyperbolic part (acoustic and material waves) of the system is stabilized globally, to take into account the coupling between variables. In the context of the reduced MHD without diamagnetic effects, the velocity profile is defined by two scalar variables: the parallel velocity ($v_\parallel$) and the velocity stream function (u)

$$\mathbf{V} = -R^2 \nabla u \times \nabla \phi + v_\parallel \mathbf{B}.$$

It is then assumed that the parallel velocity is $v_\parallel \mathbf{B}$ and the perpendicular velocity is $-R^2 \nabla u \times \nabla \phi$. In this context, we use the approximations $D_{\parallel,\mathrm{vms}}(\mathbf{b} \otimes \mathbf{b}) \simeq \frac{\delta t}{2} v_\parallel^2 \mathbf{B} \otimes \mathbf{B}$ and $D_{\perp,\mathrm{vms}}(\underline{I} - \mathbf{b} \otimes \mathbf{b}) \simeq \frac{\delta t}{2} R^4 (\nabla u \times \nabla \phi) \otimes (\nabla u \times \nabla \phi)$. The Taylor–Galerkin stabilisation associated to the density equation can finally be written as:

$$\frac{\partial \rho}{\partial t} + \mathbf{V} \cdot \nabla \rho = -\rho \nabla \cdot \mathbf{V} + \frac{\delta t}{2} \nabla \cdot \left(v_\parallel^2 \mathbf{B} \mathbf{B} \cdot \nabla \rho \right)$$
$$+ \frac{\delta t}{2} \nabla \cdot \left(R^4 \nabla u \times \nabla \phi (\nabla u \times \nabla \phi) \cdot \nabla \rho \right) + S_\rho.$$

The associated weak form now becomes:

$$\int \frac{\partial \rho}{\partial t} \rho^* = - \int \left(\nabla \cdot (\rho \mathbf{V}) - \nabla \cdot (\underline{D} \nabla \rho) - S_\rho \right) \rho^*$$
$$- \frac{\delta t}{2} \int v_\parallel^2 \mathbf{B} \cdot \nabla \rho \mathbf{B} \cdot \nabla \rho^*$$
$$- \frac{\delta t}{2} \int R^4 \left((\nabla u \times \nabla \phi) \cdot \nabla \rho \right) \left((\nabla u \times \nabla \phi) \cdot \nabla \rho^* \right).$$

Similar stabilization is applied to the other equations of the reduced MHD model. Since users perform both spatial and temporal convergence tests for production simulations, and since VMS and TG stabilization both consistently vanish in the limit of high spatial and temporal resolutions, a significant influence of the stabilization onto physics results is excluded, in particular energy and momentum conservation is not affected. Furthermore, the energy conservation diagnostics automatically evaluated during each simulation allow to check this easily. For instance in a violent MGI case, for which energy conservation is discussed in the verification section briefly, the small non-conservation observed is not associated to stabilization as separate checks confirmed. Details about numerical stabilization are omitted when discussing simulations in the results sections for that reason.

3.4. Matrix construction

The linear system of equations to be solved in each time step is described by a sparse matrix system. Due to the locality of the Bezier basis functions, usually only one out of several thousand entries in the matrix is non-zero. The matrix is constructed in a distributed way by a domain decomposition, where each MPI task is responsible for creating the matrix entries corresponding to its respective part of the finite element grid. Note that the full matrix construction could be avoided by a matrix free evaluation of matrix vector products, however earlier tests showed a higher computational cost while the memory consumption dominated by the preconditioner is not reduced too much. Each MPI task uses OpenMP to parallelize further. The matrix contributions are calculated separately for each grid element ('element matrices'). These element matrices are calculated by carrying out the integrals in the weak form of the equations over the element–local coordinates s and t via Gauss quadrature. The integration in toroidal direction is either carried out by a direct summation over equidistant toroidal planes (for very small problems), or by FFT. Vectorization and parallel scalability of the matrix construction has recently been improved [131]. Details on the matrix construction and properties are shown in reference [132].

3.5. Solver

Several variants are implemented in JOREK for solving the linear sparse matrix system after it has been assembled like described in the previous section. In the simplest approach, which is usually only used for axisymmetric $n = 0$ simulations, the whole matrix system is treated by a direct solver. Interfaces to the PaStiX [133], MUMPS [134, 135], and WSMP [136] direct sparse matrix solvers are implemented. PaStiX is mostly used in practice. PT-Scotch [137] or ParMETIS [138] are used internally by the solvers, to minimize fill-in during the matrix factorization and to achieve good performance. An interface to the STRUMPACK library [139] has been added very recently [132]. Several of the solvers offer the possibility to compress the matrix system via block low rank compression or via hierarchically semiseparable matrices, for which only first tests have been done with JOREK so far [140].

For non-axisymmetric simulations, a restarted GMRES variant [141] is usually applied [41]. The necessary linear operations on the sparse matrix system are implemented directly in JOREK on the distributed representation of the system of equations. Since the system is very stiff, good preconditioning is mandatory. For this purpose, a physics based approach is used. Since each linear eigenfunction in a tokamak is associated to a single toroidal mode number, the preconditioner assumes approximate decoupling of the toroidal harmonics. Consequently, if the matrix is written into blocks corresponding to the toroidal harmonics, the diagonal blocks describing the interaction of each harmonic with itself are kept while all off-diagonal blocks are dropped in the preconditioning. This block-diagonal structure of the preconditioning matrix allows to treat each block independently of the others. For that pur-

pose, parts of the global matrix, which has been assembled in a distributed way by domain decomposition (see section 3.4) needs to be re-distributed by an all to all communication. The complete distributed matrix remains in memory as it is needed for GMRES operations. Since this re-distribution can become inefficient for large problem sizes, an option for re-calculation of the preconditioning matrices has recently been implemented using the same routines applied for the global matrix construction [132].

Each block matrix is solved by one or several MPI tasks on one or several compute nodes, with the number of cores inside a compute node being exploited via the built-in thread support of the sparse matrix libraries (pthreads in case of PaStiX). Note that the analysis step of the solvers is performed only once in the beginning of a simulation (or upon a restart). The LU factorization of the block matrices is only performed in the first time step of a simulation, and again, when the GMRES convergence has deteriorated too much, i.e., the number of GMRES iterations in a time step exceeds a user defined threshold. Consequently, the factorized preconditioning matrices may be re-used in several or even many consecutive time steps before an update is needed.

The solve step, on the other hand has to be carried out in each GMRES iteration of every time step. In the linear phase of a simulation, the preconditioning matrix is an extremely good approximation of the complete system such that an update is not needed often and the factorization has almost no impact onto the overall performance. However, in a highly non-linear state, the factorization may need to be updated almost every time step and dominates the overall computational costs of the solver in that case. The preconditioner in such a highly non-linear state is not approximating the complete system well any more, since the energy exchange between the toroidal harmonics becomes strong. For such situations, a recently implemented generalization of the preconditioner [132] can be of advantage, where each block matrix in the preconditioner can contain several toroidal harmonics. While this improves the approximation of the complete system by the preconditioner, it also increases the computational costs for inverting the block matrices such that the best configuration for the preconditioner is case specific. In case these 'mode groups' are not mutually exclusive, several options exist for obtaining the combined result vector. Another recently implemented extension [132] allows to use complex matrix solvers for the preconditioner, reducing memory consumption and computational costs considerably.

When PaStiX is used, JOREK exploits a specific feature of the solver, where the connectivity graph is not considered separately for each dof. Instead, square blocks of dimension n_{var} are treated as dense blocks, where the number of variables typically is around $n_{var} = 8$. Exploiting this structure, greatly enhances the performance of the analysis and factorization steps and can even improve the performance of the solve step considerably. Since the result vector of the matrix system contains the *changes* of all dofs for all physical variables, the actual values of the variables need to be updated after the solve step. However, also the changes of the variables in the previous time step need to be stored, since those

are required for some of the time stepping schemes (e.g., BDF2 Gears, see section 3.2 for more information on the time stepping).

Main limitations of the solver strategy are the high memory consumption associated to the factorized preconditioning block matrices (partly mitigated by the recently added complex matrix treatment in the preconditioner), the limited parallel scalability of the direct solver used in the preconditioning (addressed by adapting the code for better optimized or further improved solver libraries), and the deteriorating efficiency of the preconditioning in case of strong toroidal mode coupling (mitigated by the recently implemented generalization of the preconditioner to integrate the most important non-linear interactions in the approximation).

3.6. Parallelization

JOREK uses a hybrid MPI plus OpenMP parallelization. The number of MPI tasks must be a multiple of number of toroidal harmonics (cosine and sine components are not counted separately here). Typically, a few MPI tasks are started per compute node and the number of OpenMP threads is adapted such that all CPU cores are exploited. The number of compute nodes used is primarily determined by the memory consumption of the simulation. As shown in reference [132], already the matrix system itself requires significant storage and the LU factorization used in the preconditioner (see section 3.5) increases this memory consumption further. JOREK is mostly used on conventional architectures with Intel or AMD CPUs including Intel Xeon Phi. Exploiting GPUs is considered for the future, in particular for the kinetic particles (see section 2.10). Base information about the simulation like the input parameters, but also the plasma state, i.e., all degrees of freedom of all physical variables on the whole grid, is duplicated across all MPI tasks (domain cloning). This duplication is affordable in terms of memory consumption, and allows an efficient load balancing for instance for the kinetic particles. The communication overhead typically associated with PiC methods using a domain decomposition is avoided since particles do not need to be transferred from one MPI task to another. An example for the scaling of the kinetic particle model is shown in figure 8.

MPI rank 0 is responsible for reading the namelist input file, broadcasting this input to the other MPI ranks, reading and writing restart files[38], printing information to the log file, etc. The grid construction and equilibrium calculation (section 2.2) is only carried out on one MPI rank typically as this is not an expensive part of the simulation. This is executed in a separate 'job' from the time evolution such that no CPUs are idling. For (large) free boundary equilibrium calculations, several MPI ranks may be necessary due to memory requirements.

The parallelization of the matrix construction during the time stepping is done by a domain decomposition such that

Figure 8. Strong scaling of a JOREK particle simulation (with 10^9 ions) from an entirely kinetic electro-static ITG turbulence simulation with JOREK on the Marconi-Fusion supercomputer (see sections 2.14 and 7.1).

the system of equations is assembled in a distributed way (see section 3.4). For the iterative solver (see section 3.5), an all to all MPI communication is carried out to extract the preconditioning matrices from the complete system (see section 3.5 for details on the solver and preconditioner). Optionally the harmonic matrices can be recomputed bringing performance gains in specific limits [132]. Handling the preconditioning matrices may be slightly unbalanced since the $n = 0$ block is smaller than the others (see section 3.4). However, MPI rank 0, which is responsible for the $n = 0$ block, either alone or together with a few other MPI tasks, usually anyway has other responsibilities, which can overlap with the solve on the other MPI ranks, such that this is not affecting performance in a significant way. The 'mode groups' described in the previous section can lead to a larger imbalance, which can be compensated by adapting the number of MPI tasks assigned to each preconditioner block matrix, which is possible in a flexible way [132]. The sparse matrix libraries applied for solving the preconditioning matrices employ a hybrid parallelization internally as well. In case of PaStiX, this is done by an MPI plus pthread parallelization, for instance. The linear operations inside GMRES, e.g., sparse matrix vector products, are well balanced, since they are based on the evenly distributed global matrix.

An example of the scalability of JOREK is shown in figure 9 for a moderate problem size. Further work to improve the performance and parallel scalability of JOREK is ongoing, as shown in section 3.8. The scalability of STARWALL is shown in figure 10, refer to reference [99] for more information about the parallelization of STARWALL and the JOREK-STARWALL coupling terms.

3.7. Code management

JOREK is written in Fortran 90/95 using Fortran 2003 and 2008 extensions in particular in the kinetic particles part. A few smaller core routines and interfaces are written in C and

---

[38] For reading/writing large files in the free boundary extension (section 2.9) and the kinetic particles framework (section 2.10), parallel I/O (MPI I/O and parallel HDF5, respectively) is used to achieve good performance. In case of the free boundary extension, this concerns reading the vacuum response matrices [99]. In case of the particles framework, it allows to avoid unnecessary communication when reading/writing large amounts of particle data.

Figure 9. Strong scaling test of the JOREK parallel efficiency for a moderate problem size on the Marconi-Fusion supercomputer using Intel Skylake CPUs. The setup is a JT-60SA ELM simulation (reference [4]) with a moderate 2D grid resolution of about 5k grid nodes and a relatively high toroidal resolution including $n = 1 \ldots 15$ included. For a fixed problem size, the number of compute nodes is increased from 16 (576 cores) to 160 (5760 cores). Note, that only 36 out of the 48 cores per node were used in this test. The relative efficiency is approximately 50% when comparing the performance between 16 and 160 compute nodes, which constitutes an excellent value for a fully implicit code. In practice, the problem size considered here would not be addressed using more than 32 compute nodes, for which the relative efficiency is around 85%.
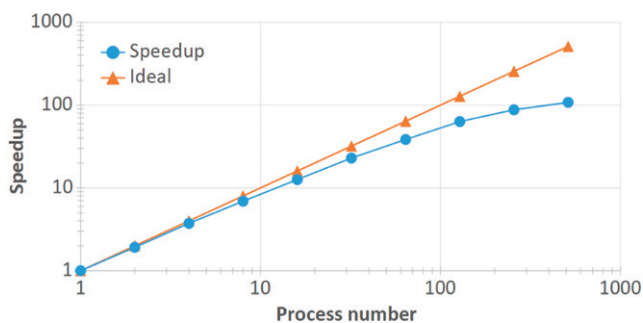
Figure 10. Strong scaling test of the STARWALL parallel efficiency for a small problem size on the Marconi-Fusion supercomputer [99] with 35 200 wall triangles equivalent to 17 600 degrees of freedom (size chosen to fit into the memory of a single compute node). The execution time of the whole STARWALL run is given. When the number of MPI tasks is increased from 1 to 128, a speed-up of 63.3 can be obtained (parallel efficiency of 50%). At 64 MPI tasks which would be used in practice, the parallel efficiency is around 60%. Larger production cases scale to a higher number of cores with a better efficiency. Note that this figure shows the scaling for response matrix calculation in STARWALL, not for the time stepping of the JOREK simulation. Calculating the STARWALL contribution to the boundary conditions in the JOREK time stepping is usually not altering the JOREK scaling significantly as this constitutes a sub-dominant and well parallelized [99] part of the computational costs except for exotic setups which combine a low plasma resolution with a very high wall resolution.

C++. JOREK presently consists of more than 250 thousand lines of source code, to which dozens of developers have contributed over the years. More than 100 pull requests are merged per year at the moment with a strongly increasing trend. More than 40 scientists presently use the code world-wide for very

different purposes and a significant fraction of those is actively contributing to the development. Via modern code development techniques and a common shared code repository, a unique stable code basis is maintained shared by all users and for all code applications (the code is hosted on the Atlassian Bitbucket system operated by the ITER Organization). A dedicated Wiki contains the collaboratively written code documentation.

A particular challenge is posed by the contradictory goals of code stability and agile development. With the strongly increasing number of users and developers over the last years, it became critical to ensure that developments for specific purposes cannot break entirely different code applications. For guaranteeing this, limited manpower was available such that a high level of efficiency was required. To achieve these goals to the best possible extent, the modular structure of the code was continuously improved and the development work flow was switched from subversion to a git based repository more than five years ago. The development work flow starts by describing issues or planned new features in the Jira issue tracker, discussing about the planned implementation, creating a branch for the respective development based on the latest code version from the main development branch, implementing the changes into that branch, raising a pull request with several code-reviewers, refining the solution iteratively, and merging the modifications into the main development branch. To ensure coordination and inform users and developers, all changes to the code are discussed and communicated in regular development meetings. At each meeting, a release version of the code is created. Smaller developments are usually merged quickly, large developments are sometimes carried out in a branch over months or even years. Measures are taken to reduce the number of long-lived branches and the remaining ones are actively kept up to date with the main development branch by regularly merging in all recent changes.

To avoid introducing problems by the numerous smaller and larger developments which take place in parallel with the large number of different physics models and code applications, we introduced a framework for carrying out automatic test cases in addition to the reviewing of pull requests. Our approach to these tests is described in the following (reference [142] explains an older approach, which has been replaced several years ago). About 50 automatic regression tests are in place at present, which are executed automatically whenever new modifications are pushed to any branch on the central repository. These test cases can also easily be executed locally by each user. The framework is flexible enough to allow for different 'versions' of the same regression test in different branches, a feature needed in case a bug is detected and resolved, or a modification is implemented that intentionally changes the code behaviour (e.g., improved grid construction). The data required for carrying out a regression test and checking the result is stored on a separate server and the correct version of the test is identified via hash keys. All regression tests are carried out with a large set of compiler/linker debugging options and some regression tests are using a second compiler (gFortran instead of Intel Fortran) to catch as many problems as possible.

The majority of the test cases are based on low-resolution non-linear simulations, which are restarted from an HDF5 restart file and continued for a single time step only. The result is then compared to a reference result with a specific absolute tolerance. This has been tested extensively to capture problems reliably and the execution in the non-linear phase ensures that all physics terms influence the solution. Among the test cases, also a few specific ones exist that work differently. The simplest ones only compile the code and all diagnostic binaries for the various physics models, other tests cover the grid construction and equilibrium solver. Specifically for the kinetic particles module, also unit tests are in place.

In spite of the considerable number of test cases, not all applications of the code can be covered, however new regression tests are continuously added and are introduced for new code features already before these are merged into the main development branch. Furthermore, a generalization of the framework is foreseen to simplify testing all combinations of different use cases and model extensions.

3.8. Outlook

This section gives a brief outlook onto developments regarding the numerical methods of the JOREK code.

For future simulations of non-axisymmetric configurations (stellarators), in addition to the derivation of a hierarchy of MHD models (section 2.15), the numerical methods also need to be adapted. In particular, the following assumptions need to be relaxed:

- The simulation grid is not axisymmetric any more, i.e. R and Z are not expressed in s and t only, but also expanded in a toroidal Fourier series. An import of such a 3D configuration from GVEC [143] has already been implemented, as shown in figure 11.
- The assumption that s and t are orthogonal to ϕ will not be fulfiled any more in general (for instance in case of PEST coordinates). In addition, a $\phi = $ const surface will in general not be planar any more (for instance in case of Boozer coordinates). This is taken into account in the ongoing implementation of the new stellarator capable models.
- In a stellarator, the linear eigenfunctions of plasma instabilities cover a toroidal spectrum instead of being restricted to a single toroidal mode number. Consequently, the presently used preconditioner, which relies on the linear decoupling of toroidal harmonics, cannot be applied any more and solver developments are important. Work in this direction is moving forward, in particular see the 'mode groups' in the preconditioner described in section 3.5 which constitutes a first important step towards efficient stellarator simulations.

Regarding the solver, several developments are on the way building up onto previous work, this includes for instance an assessment for the use of GPUs, enhanced preconditioning via an integrated iterative refinement, the use of reduced order for the preconditioner, etc. These efforts are at early stages and are thus not described here in detail.
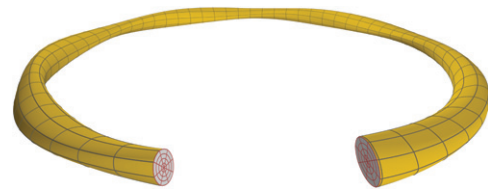
Figure 11. Example of a 3D grid (shown with very low resolution for clarity), imported from the GVEC code for a 5 field period, elliptical equilibrium, similar to the W7-A stellarator.

For the anisotropic heat transport (see also section 4), reduced order basis functions for selected quantities are being studied to further improve the numerical accuracy in particular when finite element grids are used that are not aligned to the flux surfaces, e.g., in case of VDEs.

Aiming at simplified data exchange, integrated modelling, and standardized analysis tools, the adaptation of JOREK to IMAS is on the way. The ITER Integrated Modelling & Analysis Suite (IMAS) [144] is a scientific software framework infrastructure that orchestrates execution of integrated plasma codes. The *physics data model* is a basis for coupling the plasma codes and experimental data at different space and time scales through its 60+ *interface data structures (IDS)*. A widely used IDS substructure for description of the computational domain is the *general grid description (GGD)*, with the purpose to describe any *N*-dimensional numerical grid geometry and associated plasma state quantities with time slices sampled during the simulation. Generally, JOREK couples to the equilibrium and MHD IDS for input conditions and nonlinear plasma state output that can be directly compared with other simulators that use the GGD mapping (see e.g. reference [145]). Integration of JOREK in IMAS is currently underway, with a few aspects and methods already available to allow incorporation of JOREK relevant data in the MHD IDS employing the GGD representation [146]. Spatial discretisation under GGD in 3D requires Bézier finite elements for poloidal mesh discretisation (nodes and corresponding values with derivatives) combined with real Fourier series for toroidal harmonics. Temporal slices in the MHD IDS are saved as flux aligned GGD and corresponding values on the grid. Time slices for GGD and values on GGD can be different to respect database size growth per simulation. Finally, stored simulation results enable analysis with unified tools and runs under integrated modelling workflows.

In the present grid structures, G^1 continuity is satisfied across the whole domain except for the magnetic axis and the direct vicinity of the X-point. The G^1 continuity at these special points will be restored by a special treatment with locally adapted basis functions. Furthermore, a generalization of the Bezier FE formulation to higher order basis functions and higher order continuities across element boundaries was recently implemented [147]. The implementation of an option for triangular FEs in JOREK is being considered, operator splitting or relaxation techniques are being discussed, and further refined numerical stabilization methods such as VMS and higher order Taylor Galerkin methods are being investigated.

Figure 12. Example of the G^1 continuous, i.e., continuous values and first derivatives in real space, profiles in the poloidal plane for (left) a density distribution and (right) the perpendicular conductive heat flux.

4. Verification

In this section, selected *verification* work done for the JOREK code is highlighted. This includes code–code benchmarks, comparisons to analytical solutions, comparisons between reduced and full MHD, and an assessment of energy conservation in a violent non-linear example. The emphasis here is on more recent tests, of which many have not been shown in publications and conference contributions yet. A lot more verification work can be found in the literature. E.g., the JOREK pellet model has been successfully benchmarked, for a deuterium pellet, with a dedicated code by B. Pégourié, as can be seen in figure 2 of reference [86]. Note that all the *validation* against experiments is included along with the physics studies in sections 5–7. Figure 12 shows an example to demonstrate that the finite element solutions are G^1 continuous and (if resolved well) smooth.

4.1. Convergence properties

Convergence of the exponential growth rates with the spatial grid resolution is shown in the following. Two cases are considered in circular plasmas for simplicity: a tearing mode (TM) test-case using the full-MHD model and a ballooning mode test case with the reduced-MHD model (the so-called CBM18 equilibrium for mode number $n = 20$). The grid resolution is scanned homogeneously in the radial and poloidal directions, from $(n_{\text{flux}}, n_{\text{tht}}) = (27, 180)$ to $(90, 600)$, where n_{flux} and n_{tht} are the equidistant number of radial and poloidal elements. The relative error of the growth rates scales inversely with the 5th power of the spatial resolution like $\propto (\sqrt{n_{\text{flux}}n_{\text{tht}}})^{-5}$.

With finite elements, the local error is estimated as $E \propto h^p$ according to reference [148], where h is the element size, and p is the polynomial order of the finite elements. Here, $(\sqrt{n_{\text{flux}}n_{\text{tht}}})^{-1}$ is used as an approximation of the element size h, and $p = 4$ because the Bezier elements are bi-cubic. Since the value of interest in these tests are the growth rates of the toroidal modes, which are obtained by integrating the mode energies over all the elements of the simulation domain, this adds another factor $(\sqrt{n_{\text{flux}}n_{\text{tht}}})^{-1}$ to the error estimate. Hence, the error of the growth rates is expected to scale with the 5th power of the spatial resolution, $(\sqrt{n_{\text{flux}}n_{\text{tht}}})^{-5}$.

Figures 13 and 14 show the convergence of the growth rate error, as a function of spatial resolution. For the TM, 5th order
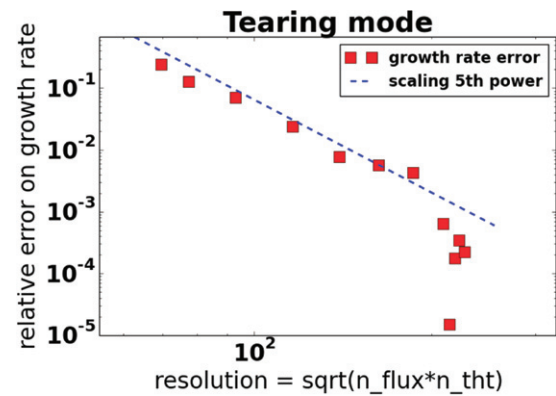
Figure 13. The relative error of the growth rate of a TM is plotted as a function of the spatial resolution $\sqrt{n_{\text{flux}}n_{\text{tht}}}$, for the full-MHD model. The error converges as the 5th power of the spatial resolution, as expected. Reprinted from [67], with the permission of AIP Publishing.
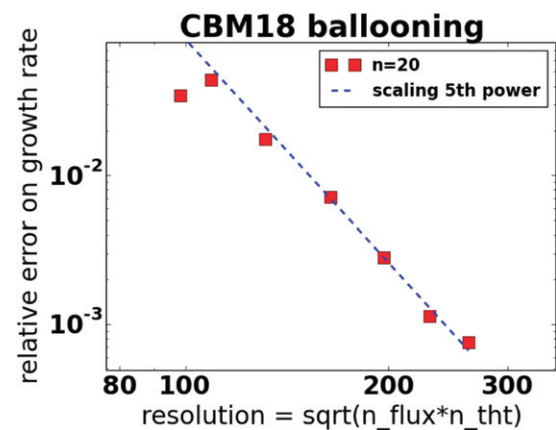
Figure 14. The relative error of the growth rate of a ballooning mode $n = 20$ is plotted as a function of the spatial resolution $\sqrt{n_{\text{flux}}n_{\text{tht}}}$, for reduced MHD. The error converges as the 5th power of the spatial resolution, as expected.

scaling is found as expected, and beyond a high enough resolution, the error diminishes dramatically, suggesting the growth rate is already fully converged. For this case, since the error diminishes faster than the expected scaling at high resolution, the intermediate case of $\sqrt{n_{\text{flux}}n_{\text{tht}}} = 2 \times 10^2$ is assumed to be the converged growth-rate. For ballooning modes, using the CBM18 case, the growth rate convergence also follows a scaling of the expected 5th order with the reduced-MHD model. For this case, since the error still diminishes at the highest resolution, the growth-rates are fitted with the expected scaling, such that the converged growth-rate is obtained by extrapolating to infinite resolution. With the full MHD model, only 3rd order scaling was found for the ballooning mode test case (see reference [149]). The reason for this different behaviour is still being investigated.

4.2. Anisotropic heat transport

The numerical treatment of anisotropic heat transport in JOREK allows to handle experimentally relevant parameters, where the diffusion coefficients along magnetic field
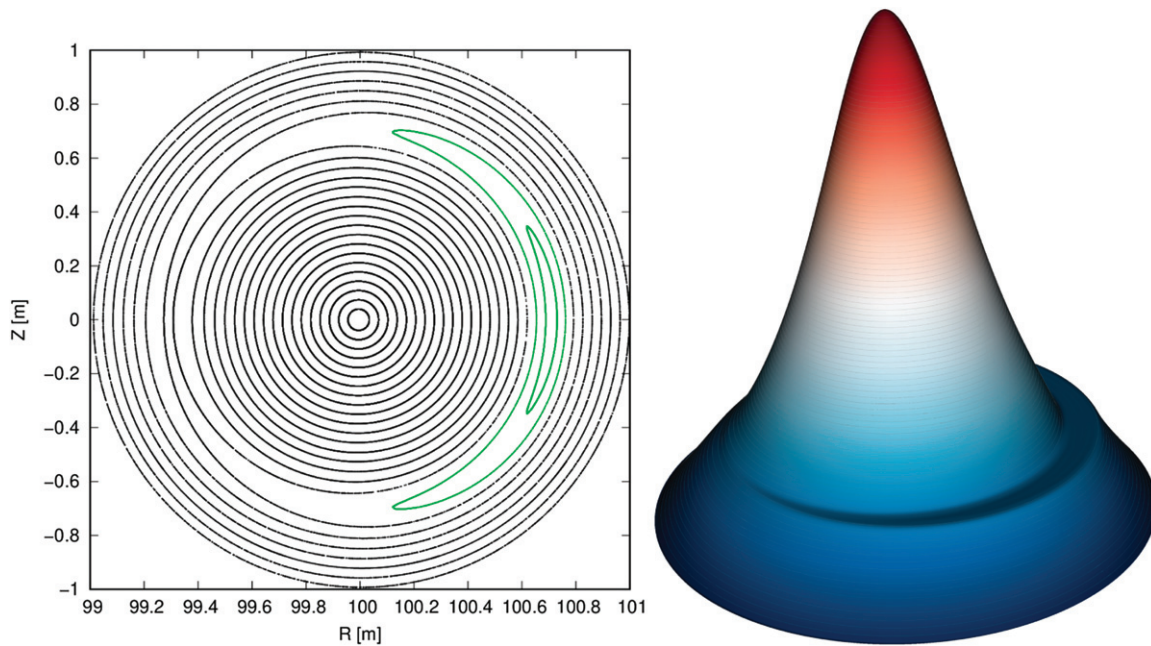
Figure 15. (Left) Poincaré plot of the magnetic field configuration used to demonstrate anisotropic heat transport in the JOREK reduced MHD model. (Right) Temperature distribution in a fully converged simulation with a heat diffusion anisotropy of 10^{10}. The temperature flattening inside the island is clearly visible.

lines and perpendicular to them differ typically by a factor 10^8 to 10^{10}. We show only a small demonstration of the properties here, based on an example with large aspect ratio circular configuration where $R = 100$ and $a = 1$. This simplified configuration is used to ease interpretation since it eliminates mode coupling and allows to initialize a single magnetic island without secondary islands or stochastization. The anisotropic heat diffusion equation is solved both using a polar grid (approximately aligned to the flux surfaces) and a rectangular grid (as example for a grid that is not flux surface aligned at all). A more detailed analysis is planned for the future in the context of implementing an advanced scheme following reference [150] for particularly demanding cases like hot VDEs with a grid not aligned at all to the magnetic field and RE fluid simulations where parallel diffusion is sometimes used as a computationally less demanding proxy for the fast parallel advection.

Figure 15 shows the magnetic configuration via a Poincaré plot as well as the temperature distribution obtained at high anisotropy. A pure 1/1 magnetic perturbation (MP) is applied to create a large magnetic island. While keeping the magnetic field fixed in time and assuming a spatially uniform density, the temperature evolution equation is solved for single large time step with the implicit Euler method to determine the steady state temperature distribution. A localized Gaussian heat source is applied in the centre of the plasma (zero inside the island surfaces). Details of the test case are described in appendix D.

Figure 16 shows cuts of the temperature distribution across the mid plane at different anisotropies and different resolutions for the polar grid (left panel). It also contains a convergence study (right panel) of the central temperature value with the grid resolution for both, aligned and unaligned grids. Note that

the error scales with the number of grid points to the third power in case of an approximately aligned grid and only with the second order in case of a non-aligned grid. When tolerating a relative error of 0.3% in the core temperature, a grid resolution of 200×200 grid points allows to easily resolve anisotropy values beyond 10^{10} in the flux surface aligned case. With the unaligned grid, the limit is around 10^9 for this resolution and error. The results would naturally change a bit with different mode numbers. However, the simulation grids tested here do not make use of any localized refinement, which is routinely applied in production simulations to concentrate the grid elements in the MHD active regions and resolve them better than the remainder of the plasma. Note that in case of a hot VDE, which corresponds to a highly anisotropic case that usually needs to be simulated in an unaligned grid, the 'spurious perpendicular heat transport' becomes negligible as soon as strong 3D MHD activity is triggered and the physical transport along stochastic field lines dominates over numerical errors.

4.3. Energy conservation

The set of equations (1)–(4) can be written in a conservative form and summed up to find the evolution of the total energy density

$$\frac{\partial e_{\text{tot}}}{\partial t} + \nabla \cdot \Gamma_{\text{tot}} = S_{\text{E}} \tag{57}$$

where the total energy density is composed by the kinetic, the thermal and the magnetic energy densities

$$e_{\text{tot}} = \frac{\rho}{2} \mathbf{V} \cdot \mathbf{V} + \frac{p}{\gamma - 1} + \frac{\mathbf{B} \cdot \mathbf{B}}{2} \tag{58}$$

and the total energy flux is
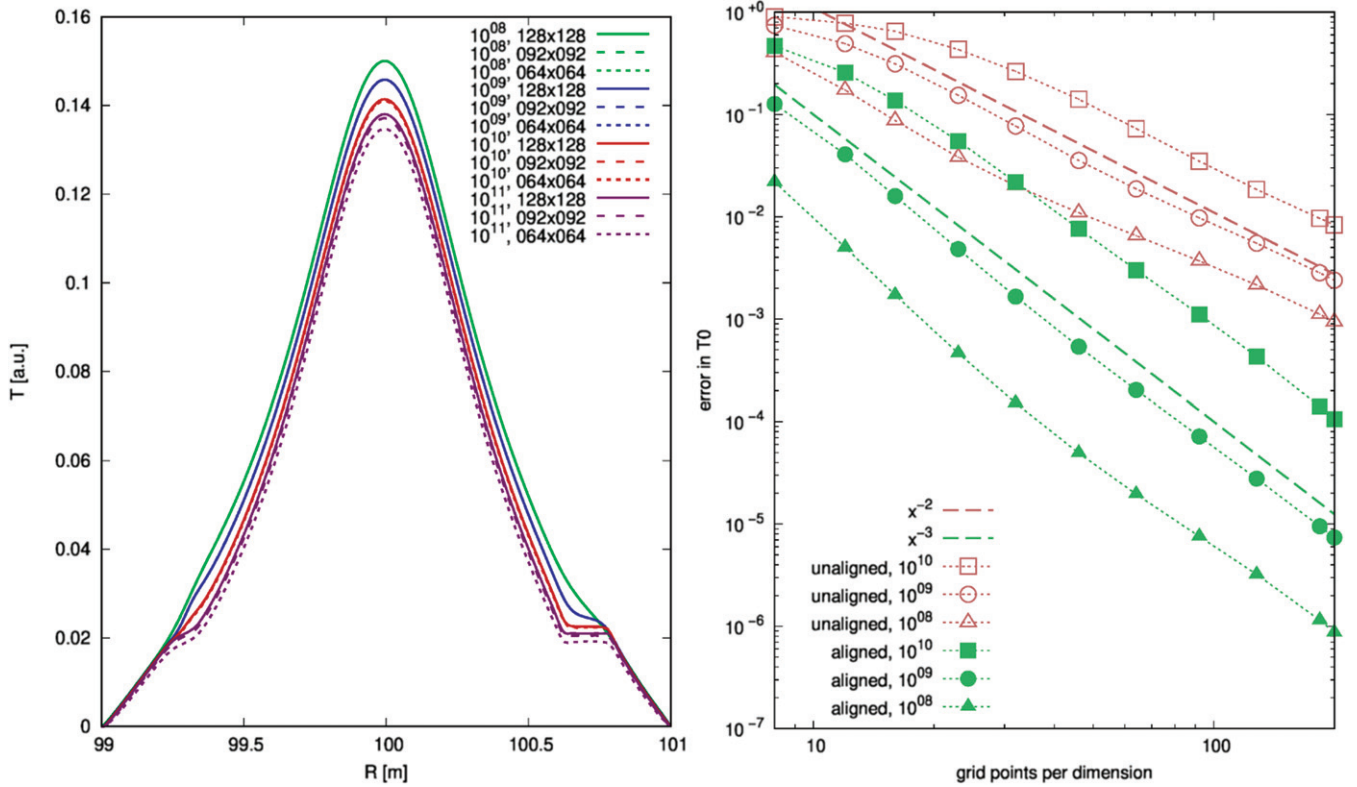
Figure 16. (Left) Temperature distribution on the midplane for the polar grid aligned to the unperturbed flux surfaces at different heat diffusion anisotropies and resolutions. The simulations shown here are done with the toroidal mode numbers $n = 0 \ldots 4$. (Right) Convergence of the relative error in the core temperature versus the number of grid points per dimension.

$$\mathbf{\Gamma}_{\text{tot}} = \left[\frac{\rho}{2} \mathbf{V} \cdot \mathbf{V} + \frac{\gamma p}{\gamma - 1} \right] \mathbf{V} - \frac{\kappa \nabla T}{\gamma - 1} + \mathbf{E} \times \mathbf{B} \quad (59)$$

equation (57) can be integrated over the plasma volume to obtain

$$\frac{\partial E}{\partial t} = \int S_{\text{E}} \, \mathrm{d}V - \oint \mathbf{\Gamma}_{\text{tot}} \cdot \mathrm{d}\mathbf{S} \quad (60)$$

where $E = \int e_{\text{tot}} \, \mathrm{d}V$ is the total plasma energy, $\int S_{\text{E}} \, \mathrm{d}V$ is the total injected power into the plasma and $\oint \mathbf{\Gamma}_{\text{tot}} \cdot \mathrm{d}\mathbf{S}$ represents the total energy flowing through the boundary (boundary fluxes). An important test for MHD codes is to check that the previous equation is satisfied also in practice, where small errors from the numerical discretization of space and time introduce errors. This can be assessed by calculating the lhs and the rhs independently. Such a test has been done for a variety of different cases. We present this comparison here for the 3D VDE case that was used for the 3 code benchmark presented in [98]. Two tests for energy conservation are shown in figure 17 (powers), where the time derivative of the total energy is compared to the sum of boundary fluxes for a 3D VDE simulation and for an RMP simulation. As the viscous and Ohmic heating terms were switched off for these particular cases, they have been calculated via postprocessing and included in $\mathrm{d}E/\mathrm{d}t$. The results indicate that energy is conserved well including the RMP case with diamagnetic drift that relies on the model with gyro-viscous cancelation. Note, that the plasma in the VDE case undergoes a very non-linear phase with stochastization of the entire domain and that all channels

for energy transport contribute to the losses (anisotropic heat conduction, Poynting flux, convection, etc). The small discrepancies at the end of the thermal quench (TQ) ($t \sim 1.2$ ms) originate from the calculation of the Ohmic heating term and are under investigation. Energy is also rather well conserved, with about 90% precision, in simulations of disruptions triggered by massive material injection (MMI) (described in section 6.2), which include Taylor–Galerkin stabilization. The slight imbalance will be investigated. The energy balance diagnostics have been implemented recently such that the various terms involved in the kinetic, thermal and magnetic energy balances are calculated automatically for any simulation, allowing an easy and systematic check of energy conservation properties.

4.4. Core instabilities

This section presents several benchmarks for core instabilities, to compare the full-MHD model against the reduced-MHD, as well as established linear MHD codes like MISHKA and CASTOR [151–154]. The first two linear benchmarks are a low-β $m = n = 1$ internal kink mode, and a low-β $m = n = 1$ tearing-mode. Both cases are run for a scan in resistivity. The kink mode is run with resistivity alone (without viscosity, and without particle or thermal diffusion), while the TM is run including all diffusive terms, with viscosity $\nu_0 = 10^{-8}$ kg m^{-1} s^{-1}, particle diffusion $D_\perp = 0.7$ m^2 s^{-1}, and heat conduction $\kappa_\perp = 1.7 \times 10^{-8}$ kg m^{-1} s^{-1}.
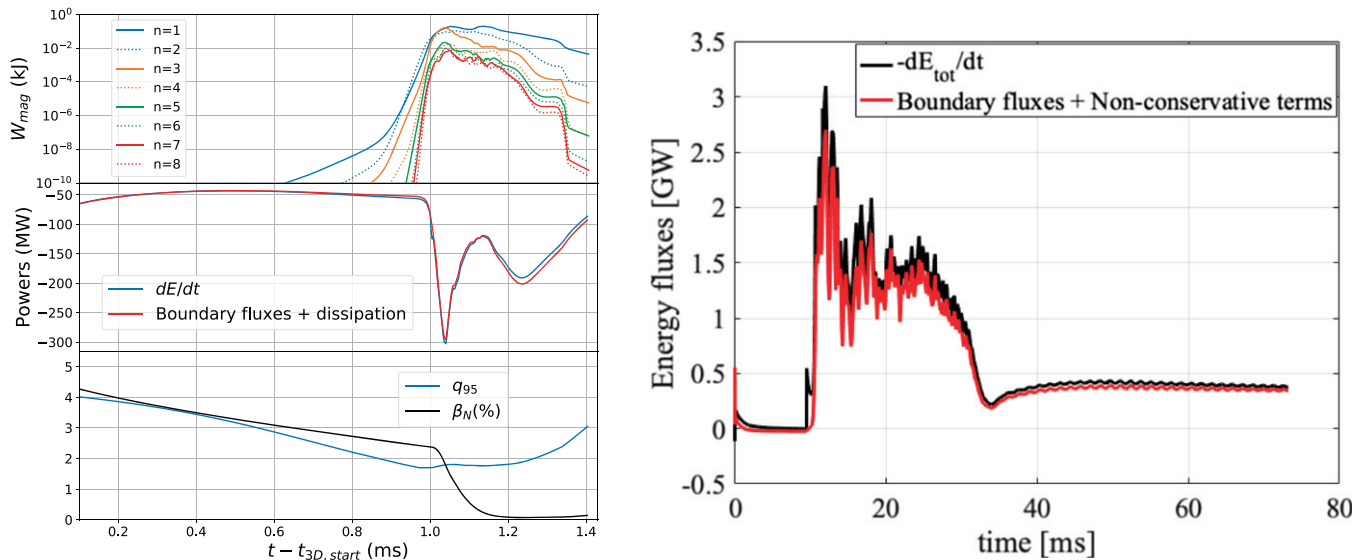
Figure 17. (Left) 3D VDE simulated with JOREK. (Top) Poloidal magnetic energy of the different toroidal harmonics. (Middle) Time derivative of the total plasma energy ($E = W_{mag} + W_{th} + W_{kin}$) and sum of dissipation powers and boundary fluxes (powers representing all energy losses). (Bottom) Evolution of the edge safety factor (q_{95}) and of β_N during the VDE. The case is explained in detail in [98]. (Right) Time derivative of the total plasma energy and the sum of non-conservative terms and boundary fluxes for a 3D RMP simulation in an ITER 15 MA plasma simulated with JOREK. The simulations includes realistic $E \times B$ and diamagnetic flows and undergoes strong MHD activity between 10 and 35 ms which affects the edge confinement. The case confirms that the error in the energy conservation introduced by gyro-viscous cancelation is small as expected from analytical estimates.

Figure 18 shows benchmarks for an internal kink mode and a TM case respectively. Poloidal cross-sections of $n = 1$ perturbed quantities are shown for the toroidal magnetic potential A_ϕ and the temperature (for the full-MHD model), and the growth rates of the modes are plotted as a function of resistivity, compared to the reduced-MHD model, and to calculations from CASTOR3D, which is also a full-MHD code [153, 154].

Although the agreement between reduced-MHD and full-MHD is reasonable for both cases, the reduced-MHD model starts to deviate from the full-MHD solution at low resistivity for the internal kink mode. This is an example where reduced-MHD becomes insufficient: internal kink modes at finite-β are not well represented like explained in reference [68]. Although this is a low $\beta_N = 0.4\%$ case, reduced-MHD already seems to be affected. At higher-β, the deviation becomes more pronounced as can be seen in the benchmark between the reduced- and full-MHD model of JOREK and the CASTOR3D code shown in figure 19. In this case, the linear growth rate of the internal kink instability is compared as a function of normalized beta. At finite beta, the reduced-MHD model clearly fails to reproduce the results of the full-MHD calculations.

The resistive layer width for a TM was simulated with JOREK visco-resistively in reference [155] and compared to analytical scaling laws as well as the Phoenix code showing excellent agreement. Furthermore, a non-linear simple TM benchmark is shown in the following. JOREK is compared here to the cylindrical full-MHD code SpeCyl [156] in simple geometry. The nonlinear verification benchmark of SpeCyl with another MHD code, PIXIE3D, is reported in reference [157]. Here, we consider the nonlinear saturation of an $m = 2$,

$n = 1$ TM in a circular tokamak with large aspect ratio in the zero-β limit. The simulations are performed at Lundquist number $S = 10^6$ in the limit of negligible viscosity. The results from the two codes are compared in figure 20. It is observed that, despite the different geometry (toroidal vs cylindrical) and models (reduced vs full-MHD), the two codes agree very well for this strong guide-field, large aspect-ratio problem. In particular, both the linear growth and the nonlinear saturation of the TM resulting from the two codes turn out to be in very good quantitative agreement (the maximum deviation of about 5% can be attributed to the difference between cylindrical and large-aspect ratio geometries).

4.5. Edge instabilities

Comparisons have been made between JOREK and BOUT++ with the aim to validate the nonlinear MHD codes. BOUT++ [7] is a framework for plasma fluid simulations, the model used for the comparison is given in reference [158] and for JOREK in section 2.3. The comparison was performed with a largely simplified geometry, a circular plasma with an aspect ratio of 3.3, $q_{95} = 3.0$, a uniform density (1.0×10^{19} m^{-3}), and a hyperbolic tangent fit is used for the temperature where $T_e + T_i = 19$ keV in the core decreasing to 94 eV at the edge of the plasma. The simulations performed have a resistivity of $6.1 \times 10^{-6} \Omega$ m and the kinematic viscosity is 117 m^2 s^{-1}. The differences between the two codes are detailed in [159] and include some deviations between the actual models (equations), evolution of equilibrium quantities, numerical methods, grids and boundary conditions. Despite some differences, the linear benchmark showed agreement within 4% between BOUT++ and JOREK with and without diamagnetic effects as seen from figure 21. Setting up a non-linear
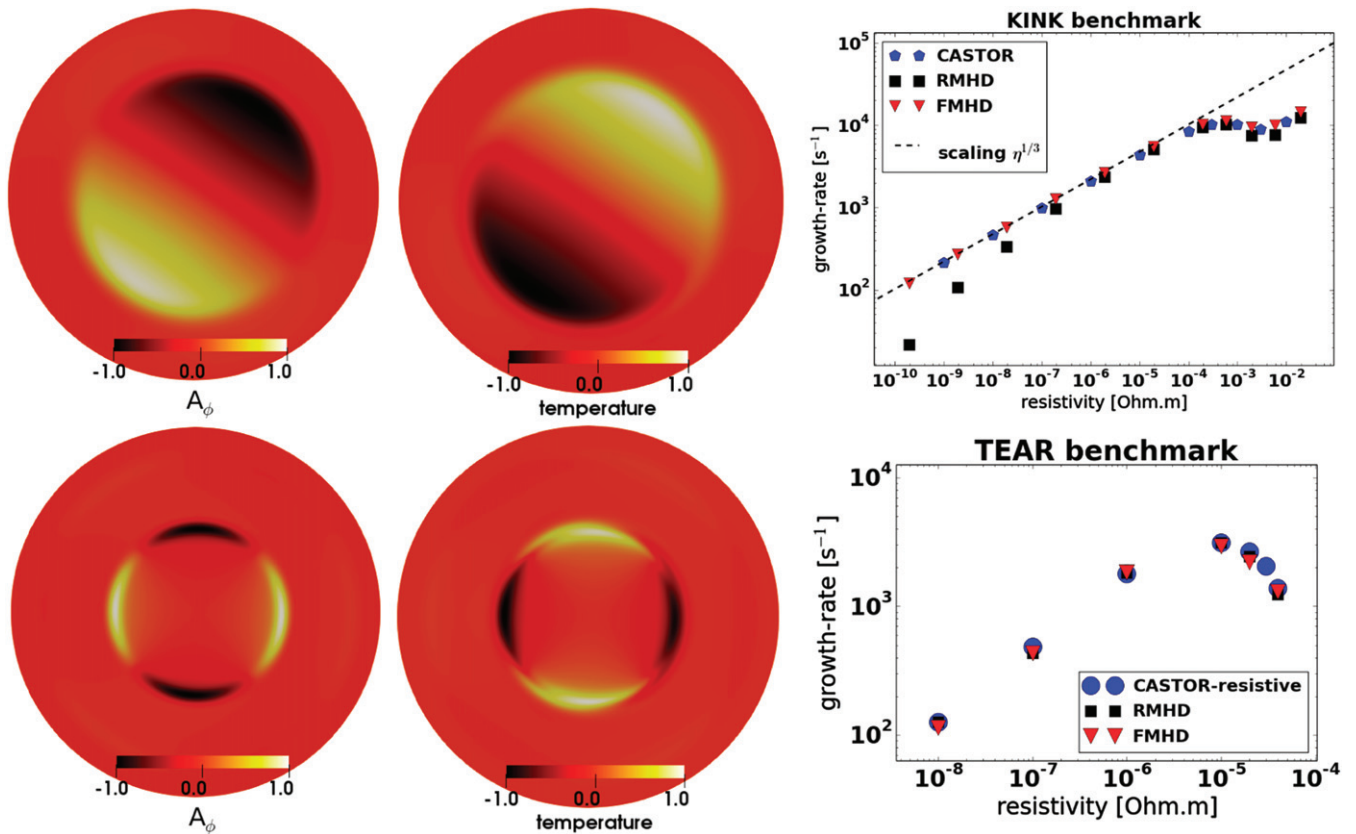
Figure 18. 1/1 kink (top) and 2/1 TM (bottom) test cases. Normalized perturbations are shown of the toroidal component of the magnetic vector potential (left) and the temperature (middle). Linear growth rates from the JOREK reduced and full MHD models are compared to results from CASTOR3D. Reprinted from [67], with the permission of AIP Publishing.
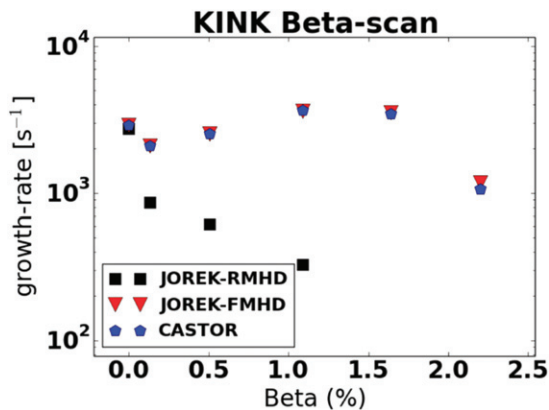
Figure 19. Linear growth rates of an internal kink instability in a plasma with circular cross section as a function of normalized beta, obtained by the reduced- and full-MHD models of JOREK and the linear full-MHD code CASTOR3D. Reprinted from [67], with the permission of AIP Publishing.

comparison based on this case would require further efforts in resolving some of the differences between the codes, e.g., regarding the boundary conditions and hasn't been attempted as of now.

A benchmark of peeling-ballooning modes (PBMs) is done using an X-point JET-like plasma. Refer to reference [65] for details regarding the benchmark configuration. Figure 22 shows the growth rate with and without diamagnetic effects, as

a function of toroidal mode number, for both reduced- and full-MHD. The agreement between the two models is reasonable with deviations below 5% (highest deviations for the largest mode numbers that are most critical in terms of resolution), particularly considering that this includes $E \times B$, parallel and diamagnetic flows in the pedestal and SOL.

The ballooning mode rotation obtained in JOREK simulations has been successfully validated against analytical linear computations in reference [160]. In the laboratory frame, the $\mathbf{E} \times \mathbf{B}$ velocity should be added to the mode velocity when comparing with experimental measurements [161].

4.6. Scrape-off layer

The SOL modelling in JOREK at this point is still fairly simplified, while several attempts are presently on their way to improve—via a better representation of low temperature physics in the fluid picture as well as via a kinetic treatment of neutral particles. In spite of the incomplete picture presently available, comparisons have been made between JOREK and SOLPS [162] to get a first assessment of the status of the JOREK SOL model. The comparison demonstrates the capabilities of the diffusive neutrals model [163], which was used to perform the simulations. Full details and results are given in [159].

A double-null H-mode MAST-U Super-X case was used in an attempt to obtain a detached divertor for ELM burn-through
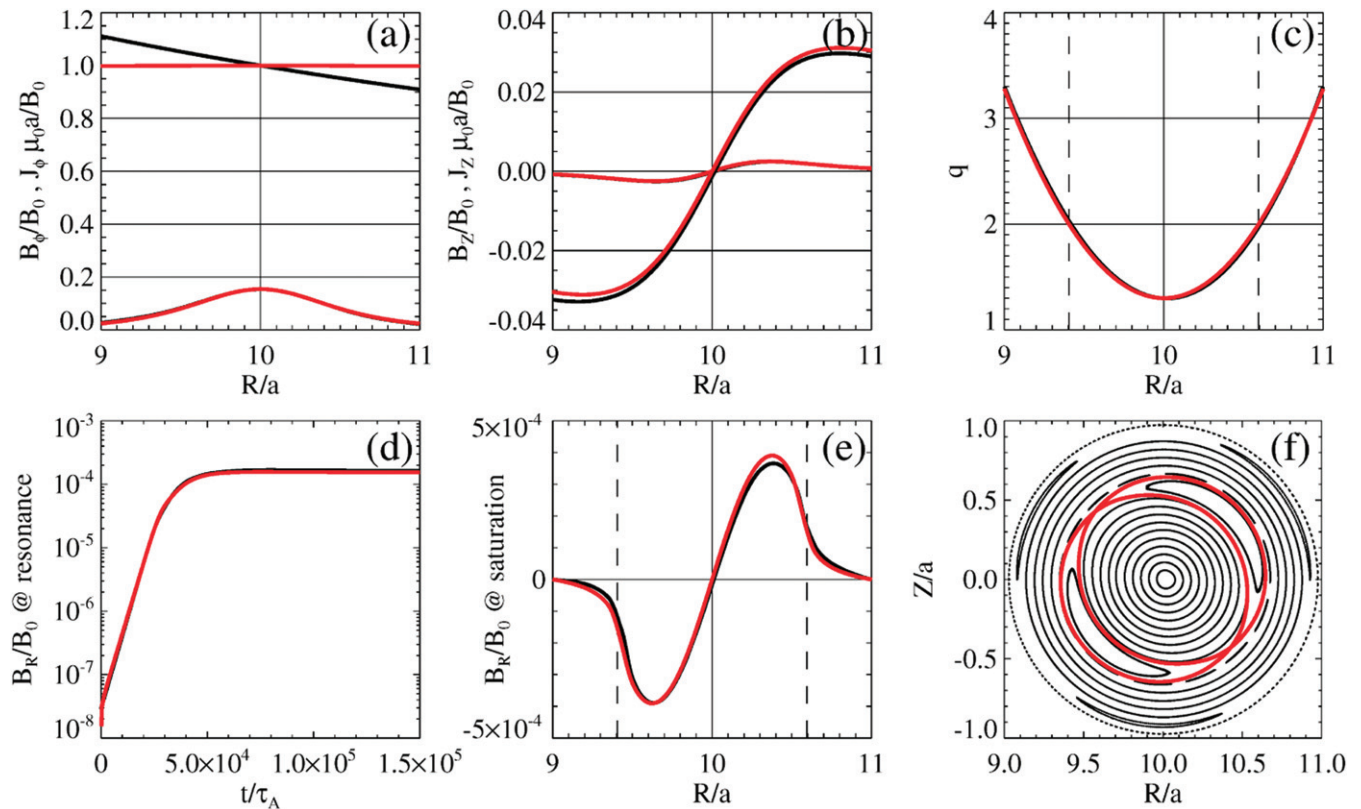
Figure 20. Nonlinear verification benchmark between JOREK and the cylindrical full-MHD code SpeCyl. A 2/1 TM is considered in a circular tokamak with large aspect ratio $R/a = 10$ in the zero-β limit. Black and red curves are for JOREK and SpeCyl calculations, respectively. Initial equilibrium profiles along the horizontal diameter are shown in the first row: (*a*) magnetic field and current components in the toroidal direction, (*b*) magnetic field and current components in the Z direction, and (*c*) safety factor. The TM nonlinear evolution is shown in the second row: (*d*) temporal evolution of B_R at the $q = 2$ rational surface, (*e*) B_R profiles along the mid plane at in the non-linearly saturated state, and (*f*) Poincaré plot at nonlinear saturation from JOREK, with the separatrix of the 2/1 saturated island from SpeCyl overplotted in red.

studies. This case had a toroidal field of 0.64 T, plasma current 1 MA, $q_{95} = 7.9$, a central density and temperature ($T_e + T_i$) of 5.2×10^{19} m^{-3} and 1.8 keV respectively. Parameter scans were performed and compared to the SOLPS results in [164] to test the JOREK diffusive neutrals model. Figure 23 shows the upstream density scan. Due to the missing physics, in particularly charge exchange, it was not possible to obtain a roll-over as steep (deeply detached plasma) as the SOLPS simulations. Nevertheless, a roll-over in the target parallel electron density flux was obtained as the upstream density increased, the target electron temperature decreased to a few electron volts and the ionisation front moved from the target—indicating a detached regime was obtained.

4.7. Runaway electron fluid

Verification of the RE fluid model in JOREK was performed with respect to (a) the conversion of thermal current into an RE beam in a simple axisymmetric plasma via comparison to a lower dimensional code, and (b) regarding the linear growth of resistive kink modes. Further efforts are on their way.

Benchmarking for the current conversion was done with the one-dimensional RE code GO [165], by triggering an artificial TQ in a large aspect ratio circular plasma ($R = 10$ m and $a = 1$ m) by imposing a large perpendicular thermal

diffusivity. The parameters at the initial equilibrium state of the plasma were $I_p = 0.67$ MA, on-axis toroidal magnetic field $B_{\phi,0} = 1$ T, central temperature $T_0 = 1.7$ keV, central density $n_0 = 1 \times 10^{20}$ m^{-3} and central resistivity $\eta_0 = 1.1 \times 10^{-7}$ Ω m. The large perpendicular diffusion leads to a drop in the core temperature of the plasma to about 25 eV in a time of about 60 ms. Thresholds were set to initiate the Drecier generation when $E_{\parallel}/E_D \geqslant 0.01$ and the avalanching when $E_{\parallel}/E_c \geqslant 1.7$. The temperature profile evolution in GO is taken as an input from JOREK. Figure 24 (left panel) shows an excellent agreement between the result obtained with JOREK and GO for the evolution of the total and RE currents (deviations around 1%). The simulations also show the often observed central peaking of RE current profile, as can be seen in the pre-quench and post-quench current density profiles shown in figure 24 (right panel). Advanced source terms that incorporate the influence of partially ionized impurities onto the RE generation [166] are presently being implemented and validated.

Verification with respect to linear growth of the resistive internal kink mode is done by considering that a certain fraction of the equilibrium plasma current is assumed to be carried by REs, wherein the RE current density has qualitatively the same profile as the total current. Both, thermal and RE
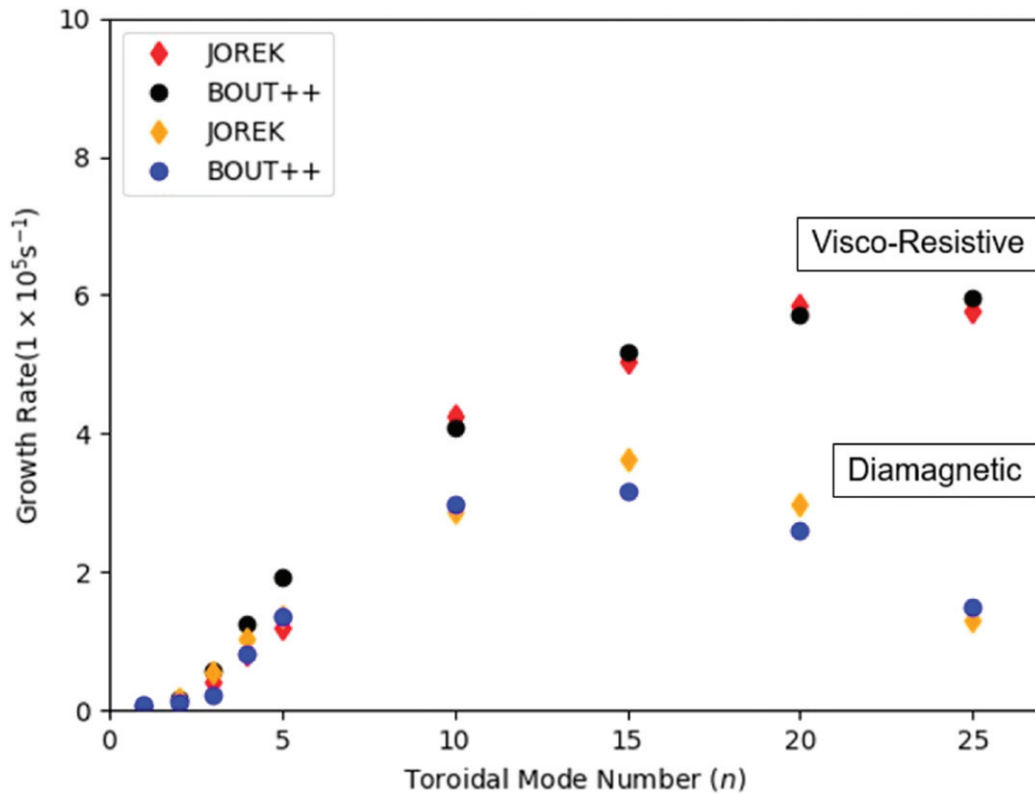
Figure 21. Linear growth rates for the different toroidal mode numbers are compared between BOUT++ and JOREK for visco-resistive simulations, and for visco-resistive simulations including diamagnetic flows.
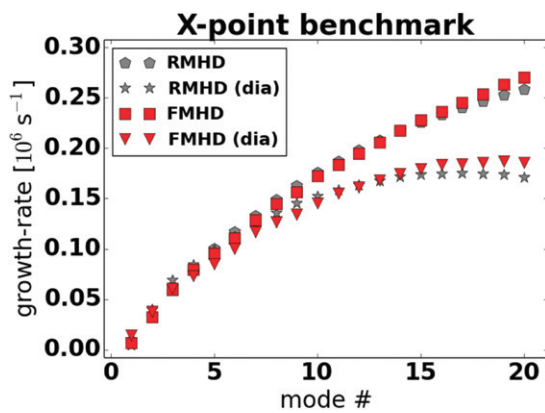
Figure 22. The growth rate of the PBMs as a function of toroidal mode number n, with and without the diamagnetic effects, for both the reduced- and full-MHD models. Reprinted from [67], with the permission of AIP Publishing.

Figure 23. The target parallel electron density flux and target electron temperature as a function of upstream density, comparing JOREK with SOLPS results from [164] for an MAST-U Super-X H-mode case. Reproduced courtesy of IAEA. Figure from [164]. © EURATOM 2020.

background current densities are kept fixed in time. A large aspect ratio circular plasma ($R = 10$ m and $a = 1$ m) in a fixed-boundary static equilibrium ($\mathbf{v} = 0$) is chosen with parameters $B_{\phi,0} = 1$ T, $I_{\mathrm{p}} = 0.31$ MA and on-axis temperature $T_0 = 48$ eV. The equilibrium is ($m = 1$, $n = 1$) kink unstable with the $q = 1$ surface within the plasma. Thermal and mass diffusivities, all the sources (including RE generation) and RE advection were set to zero, while the resistivity was assumed to be temperature independent and spatially constant. Figure 25 shows the linear growth rate of the internal kink mode as a function of normalized resistivity (inverse Lundquist number S^{-1})

for the various fractions of RE current considered. It can be observed that an increase in the RE current fraction leads to a gradual recovery of the low resistivity analytical scaling $S^{-1/3}$ even at large values of the normalized resistivities. This is primarily due to the reduced effective resistivity in the presence of REs. That is, when the RE current fraction is increased, the region outside the resistive layer tends towards the ideal MHD limit in which the low-resistivity analytical scaling is valid.
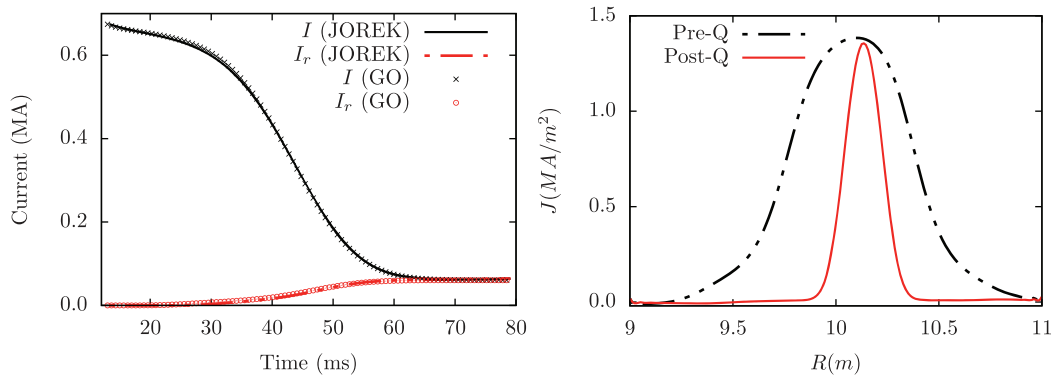
Figure 24. (Left) Time evolution of the total plasma current I and the RE current I_r during the current quench phase. (Right) Midplane current density profiles before and after the current quench obtained from JOREK, showing a relatively peaked RE current profile. Reprinted figure with permission from [111], Copyright 2019 by the American Physical Society.
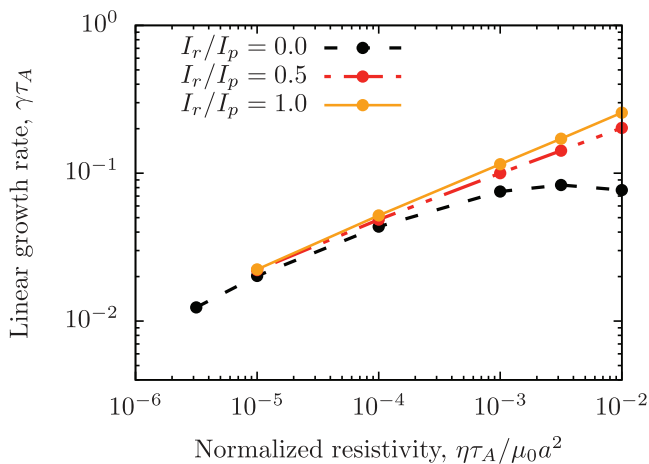
Figure 25. Linear growth rate of the resistive internal kink $(1, 1)$ mode as a function of the normalized resistivity at various initialized fractions of runaway current. Here, $\tau_A = a\sqrt{\mu_0\rho_0}/B$ is the Alfvén time and γ is the growth rate in SI units. Reprinted figure with permission from [111], Copyright 2019 by the American Physical Society.

Such a behaviour has also been observed by Matsuyama *et al* [167] in a similar (but not identical) case.

4.8. Free boundary simulations

Regarding free boundary simulations, a lot of verification work has been performed. This includes the benchmarks shown in reference [88] for a free boundary equilibrium and the linear growth rates of TMs in the presence of a conducting wall. In reference [168], the growth rate of resistive wall modes was compared to analytical values and a comparison of VDEs in a simplified ITER-like geometry was compared to the CEDRES++ code including realistic wall resistances. Basic verification tests (not shown here) have been performed including convergence checks, the induction of currents by the plasma in a conducting wall or passive coil, the mutual interaction between conductive structures and a comparison of an axisymmetric and non-axisymmetric coil field on the JOREK boundary employing the Biot–Savart law [72, 92, 169]. Further validation work is included in reference [92] and some more recent and more advanced tests are shown in the following.

A non-linear benchmark with the DINA code [92] showed good agreement for the evolution of the LCFS during an ITER 15 MA axisymmetric VDE as shown in figure 26(*a*) with differences between the shapes of the last closed flux surfaces of only few centimeters.

An international benchmark [96, 170] between the JOREK, M3D-C^1 and NIMROD codes for an axisymmetric VDE in an NSTX-like plasma has shown excellent agreement not only on the growth rate of the instability, the evolution of the radial and vertical position of the magnetic axis as well as of the toroidal plasma and wall currents, but also on the spatial structure of the halo currents flowing into the wall (see figure 26(*b*)). The agreement is remarkable in particular, since the models implemented in the codes differ significantly (e.g., full MHD versus reduced MHD; Greens functions approach versus a direct discretization of the vacuum region with conductors in the computational domain).

To test the mutual interaction between plasma, coils, and passive conductors, a benchmark was performed with the DINA code. The plasma motion resulting from a prescribed time evolution of currents flowing in coils was simulated to verify that the plasma dynamics are captured correctly in the presence of time varying coil currents [171]. For that purpose, an ITER 7.5 MA/2.65 T case was chosen where the plasma position was modified in time by the prescribed coil currents for ELM control studies. The oscillation in the vertical position is produced by applying the current waveform shown in figure 27(*a*) in the in-vessel vertical stability coils (VS coils). To get good agreement for the plasma motion (figure 27(*b*)), active and passive conductive structures had to be modelled accurately and the interaction between coils, passive structures and plasma needed to be captured correctly (see figure 44 for the setup of the conducting structures) as the time evolving coil currents induce wall currents which further modify the overall plasma motion. For this simulation, the ITER vacuum vessel was modelled as two thin stainless steel toroidal shells with a width of 6 cm each, the conducting outer triangular support and divertor inboard rail are also included.

The international three-code axisymmetric benchmark between the JOREK, M3D-C^1 and NIMROD codes shown above has been recently extended to 3D VDEs [98]. The run
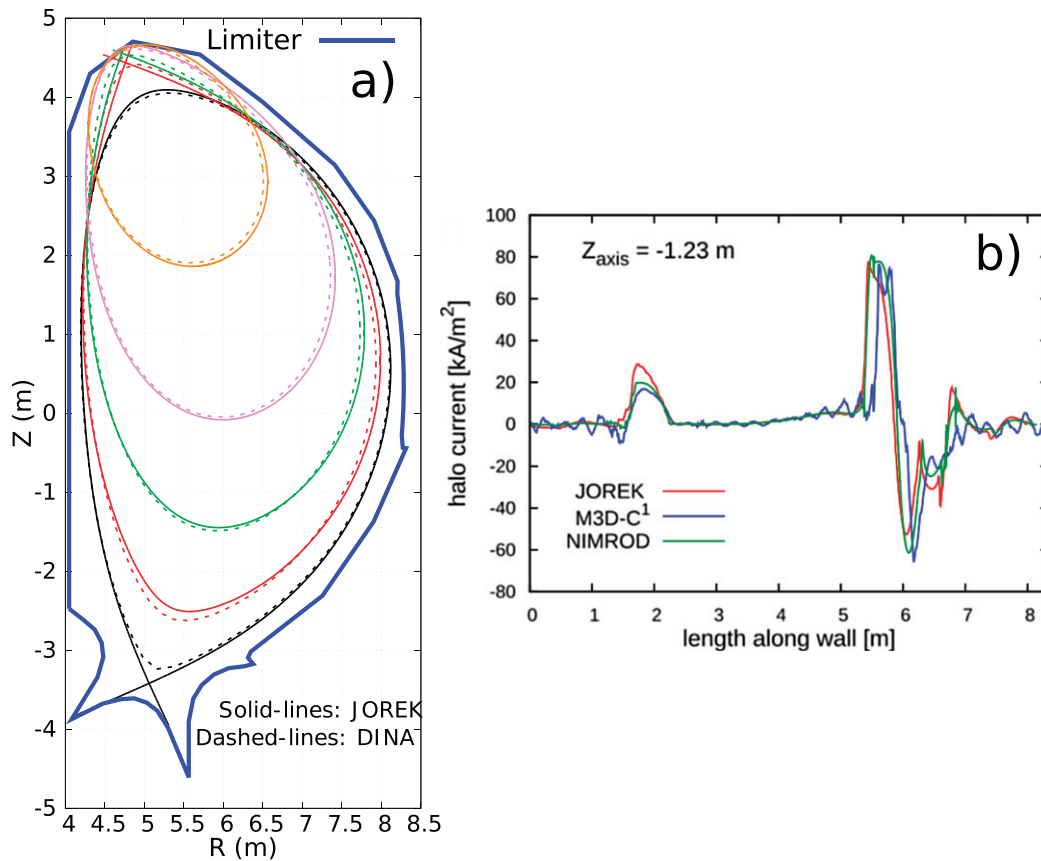
Figure 26. Non-linear 2D axisymmetric benchmarks for VDEs. (*a*) JOREK/DINA benchmark where the evolution of the LCFS is compared [92]. (*b*) JOREK/NIMROD/M3D-C^1 benchmark where the normal component of the current density is compared as a function of the distance along the wall during the late stage of a VDE. Panel (*b*) is reprinted from [97], with the permission of AIP Publishing.
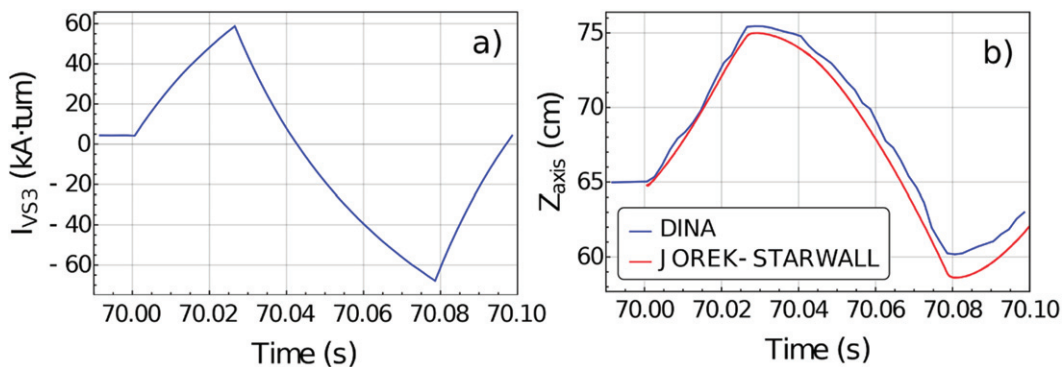

Figure 27. Non-linear 2D axisymmetric benchmark for a vertical position oscillation of an ITER 7.5 MA/2.65 T plasma. (*a*) Current waveform for the ITER vertical position control coil VS3. (*b*) Vertical position of the magnetic axis in DINA and in JOREK. Panel (*b*) is reproduced courtesy of IAEA. Figure from [172]. © 2018, ITER Organization.

was divided into two phases: an axisymmetric part for the early evolution of the VDE and a 3D simulation for the MHD active phase. The 3D run was started when the plasma became limited by the wall instead of the lower X-point. The toroidal harmonics that have been included in JOREK and NIMROD for the 3D phase are $n \in [0, 10]$ and M3D-C^1 discretized the toroidal direction with 16 Hermite cubic elements. The JOREK simulation was run with the reduced MHD model and the other codes employed full MHD models. In spite of pronounced differences on models for plasma and wall as well as

numerical methods, the three-dimensional features are in very good qualitative agreement. The three codes show that the plasma is unstable to low-*n* external kink modes 0.85–1.1 ms after the plasma becomes limited by the wall with a dominant $n = 1$ component of similar amplitude in all codes. This happens when the $q = 2$ rational surface moves into the open field-line region (see figure 17). The development of the kink modes causes a stochastization of the magnetic field lines triggering a TQ with a duration of 0.14 ms in JOREK, 0.24 ms in M3D-C^1 and 0.18 ms in NIMROD. The evolution of the
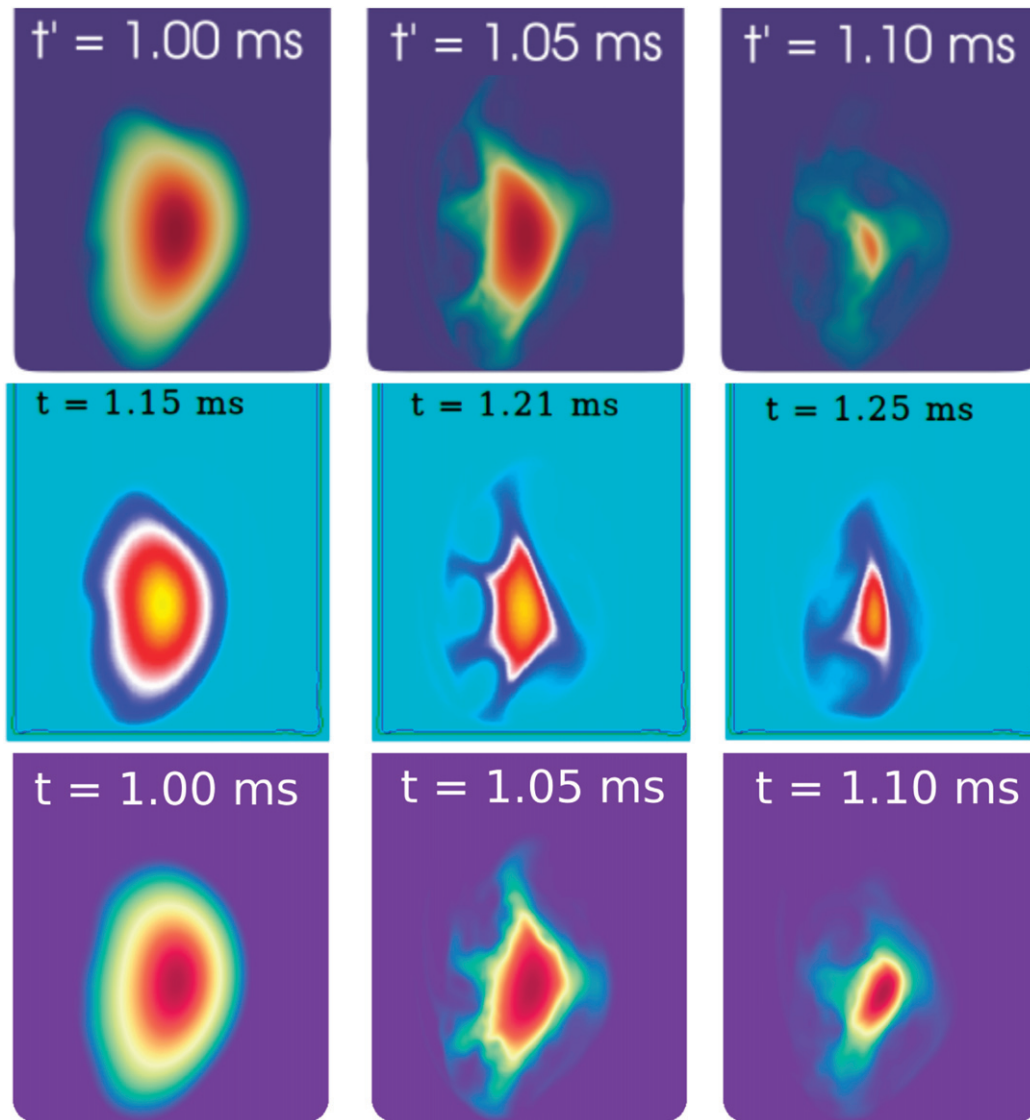
Figure 28. Evolution of the pressure in the $\phi = 0$ plane over time as computed by JOREK (top), M3D-C^1 (middle) and NIMROD (bottom) in arbitrary units. Re-print from reference [98].

thermal pressure during the TQ is shown in figure 28, revealing similar filamentary structures for the three codes. Figure 29 shows the evolution of the magnetic energy of the $n = 1$ harmonic and the vertical and horizontal wall forces in the three codes. The vertical force is in good agreement between the codes throughout the simulations except for a small shift in time and the more sensitive horizontal force is in reasonable quantitative agreement, with NIMROD showing a factor two smaller force than the other two codes. In spite of some differences, the benchmark demonstrates that the very different numerical descriptions of the resistive wall structures used in the codes lead to comparable results for such a violent 3D VDE case and that the ansatz based reduced MHD model used here for the JOREK simulation is capturing the 3D dynamics of the wall currents, even for the large β spherical plasma considered here. The full MHD models of M3D-C^1 and NIMROD found maximum toroidal asymmetries in the plasma current (I_p) of 1%–4%. However such

asymmetries were not found in the JOREK simulations due to the employed BC for the magnetic field [88].

A benchmark with CASTOR3D was carried out to test the effect of passive coils on the plasma in an ASDEX Upgrade case where the main passive stabilization against vertical displacement is due to the PSL. The PSL conductivity was varied to compare the dependence of the VDE growth rate on this conductivity. In the setup, no conducting wall was included for simplicity to only assess the coil effects. A high plasma resistivity in the SOL region was required to get agreement with CASTOR3D, which treats the SOL as part of the vacuum. The results in figure 30 show good agreement in the linear growth rates over many orders of magnitude in the PSL conductivity with largest deviations around 20% in the limit of high PSL conductivity where small differences in numerical treatment of the PSL geometry are likely dominating the errors. This includes the no wall limit, where the growth rate is not
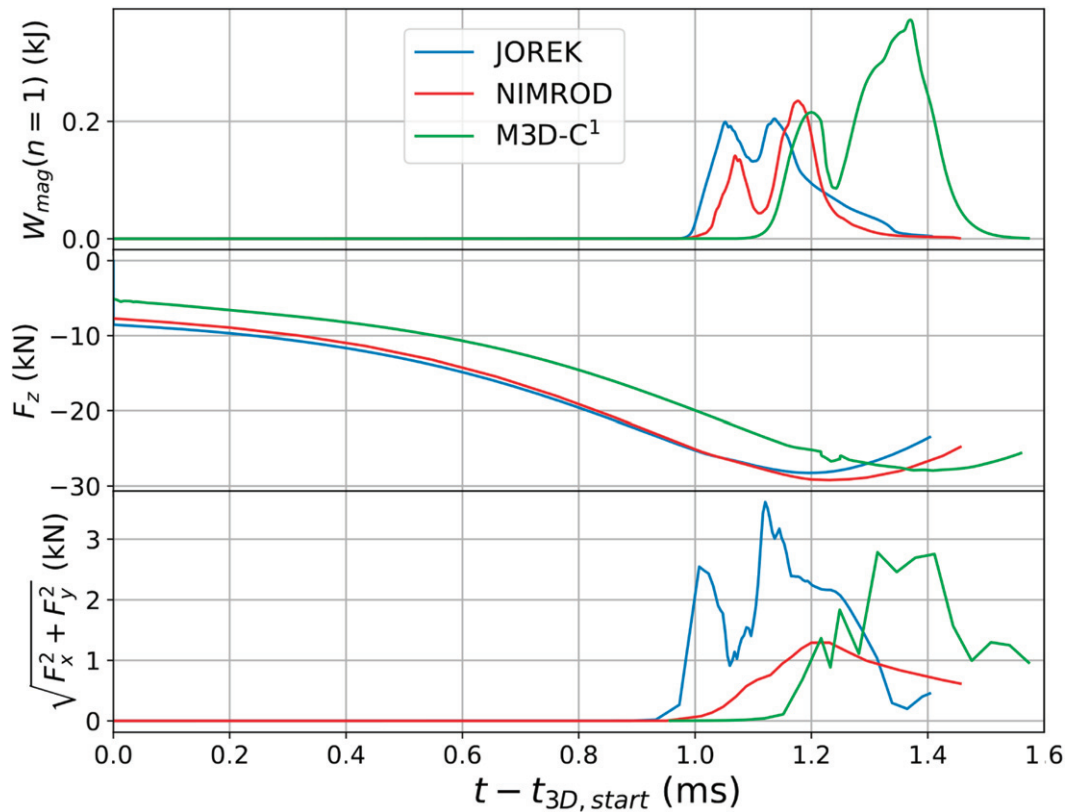
Figure 29. 3D VDE benchmark between JOREK, NIMROD and M3D-C[1]. The case and comparison is explained in detail in reference [98]. (Top) Magnetic energy of the dominant mode ($n = 1$). (Middle) Total vertical force. (Bottom) Total horizontal wall force in JOREK, NIMROD and M3D-C[1]. Re-print from reference [98].

depending on the PSL conductivity any more, but on the plasma inertia.

4.9. Kinetic particles

For non-relativistic particles, solvers are available for both full orbit and guiding respectively gyro centres. The well-known Boris method is used for the full orbit following. The implementation of the Boris method has been verified in detail in reference [172]. For guiding centres, both the standard Runge–Kutta (RK4) and the variational method by Qin [173] in the modified form [174] have been implemented and verified. For the gyro-center following, RK4 with an orbit averaged electric field is used. The solvers have been verified following a banana orbit for 26 periods in a static tokamak equilibrium without electric field. This standard benchmark case has been described in detail in references [173, 174]. The error is defined as the time averaged difference with the initial values. The scalings follow the expected order as a function of the time step. The Boris method conserves energy to a level close to machine precision. The error in the toroidal momentum scales quadratically with the time step. The RK4 method, being a fourth order method, should yield an error scaling with the 5th power of the time step, but integrated over a given time interval the error is expected to scale as the 4th power. Figure 31 shows a scaling between 4th and 5th order for the error in both the toroidal momentum as well as the energy. The variational method is expected to conserve the toroidal momentum up to

machine precision. The observed error actually decreases with increasing time step. The error in the energy scales quadratically up to a time step of 1000 gyro-periods. For larger time steps the modified Qin method is unstable (in agreement with reference [174]). As the amount of work (i.e. computing time) is about equal for the RK4 and Qin methods, it appears that RK4 is significantly better for all values of the time step. However the variational method has the advantage that the error is bound in time, i.e. the error will not increase when the integration time is increased. This is in contrast with the error in the RK4 method which tends to increase with time. In this particular case, the error of the variational integrator becomes smaller than the RK4 error for an integration time of the order of 0.1 s.

4.10. Relativistic kinetic particles

The relativistic particle tracing in JOREK has so far mainly been used to study RE confinement during plasma disruptions (section 6.4). Similar studies have been done with an orbit-following code ASCOT5 [175, 176]. In both codes, full gyro-orbits are solved with the volume-preserving algorithm [177] and the (fixed time-step) guiding centre orbit is solved with RK4. The main difference between ASCOT5 and JOREK-particles is that ASCOT5 does not use finite elements for the calculation of fields; the magnetic field is interpolated in a uniform cylindrical grid with cubic splines instead. Because both codes are being used to study transport of REs in a perturbed
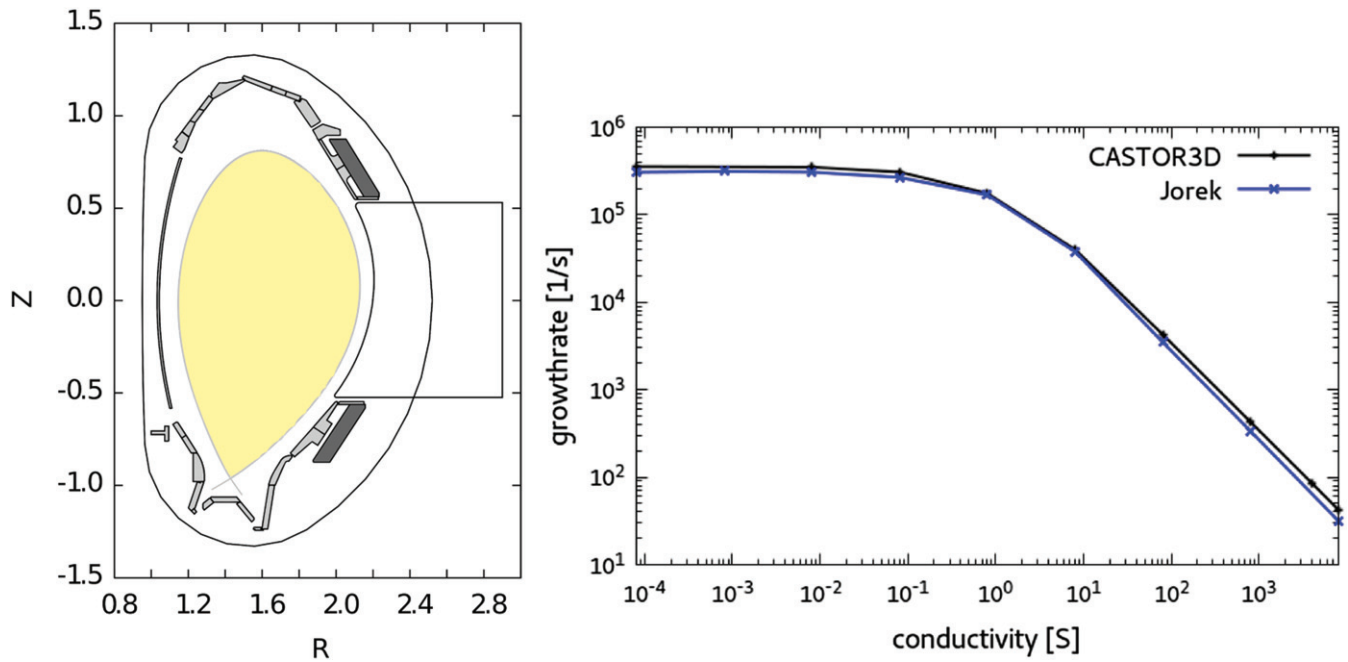
Figure 30. The figure on the left shows the conductive structures in ASDEX Upgrade including the passive stabilisation loop (PSL) (massive structures above and below the sketched port marked in dark grey). As the conducting vacuum vessel is located at a large distance from the plasma, the PSL is needed to reduce the growth rate of the vertical motion. On the right, the growth rate for different PSL conductivities is shown. For a low conductivity, the growth rate is limited by the inertia of the plasma as the coil resistance is too high to allow for significant currents. For large conductivities, the PSL resistance determines the growth rate.
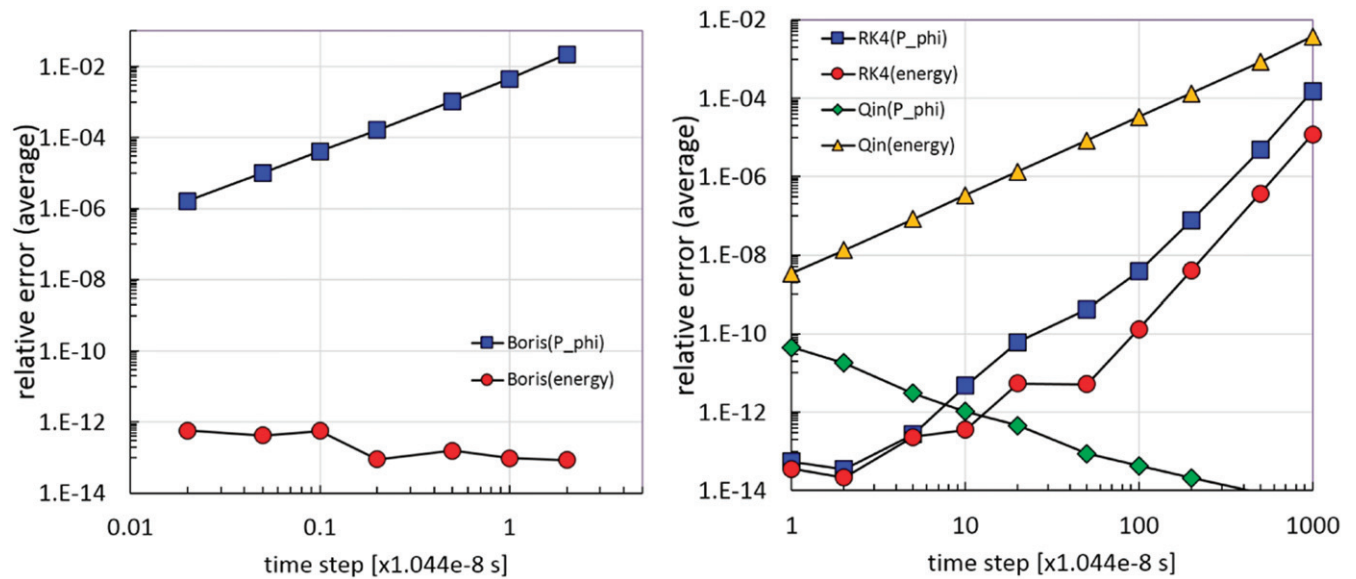
Figure 31. Scaling of the relative error on energy and toroidal momentum for the Boris, RK4 and Qin solvers for single banana particle orbit (as defined in references [173, 174]). The time step is normalised to the ion cyclotron frequency. The total integration time is fixed at 10^6 (0.01 s).

magnetic field, it is reasonable to verify that the codes yield equivalent results.

The TQ phase of a JET disruption was chosen as test case, where the stochastic field line region extended all the way from the edge to the core. Both electric and magnetic fields were included and the fields where set to be stationary in time. JOREK postprocessing tools were used to evaluate the electro-

magnetic field on a cylindrical grid and the data was exported to ASCOT5. It was verified that the grid was dense enough as not to affect the benchmark results significantly by repeating the simulations with different grid resolutions.

For the benchmark, a marker population of 5000 passing 1 MeV electrons was initialized on the outer mid plane on a fixed radial position. The markers were spread toroidally
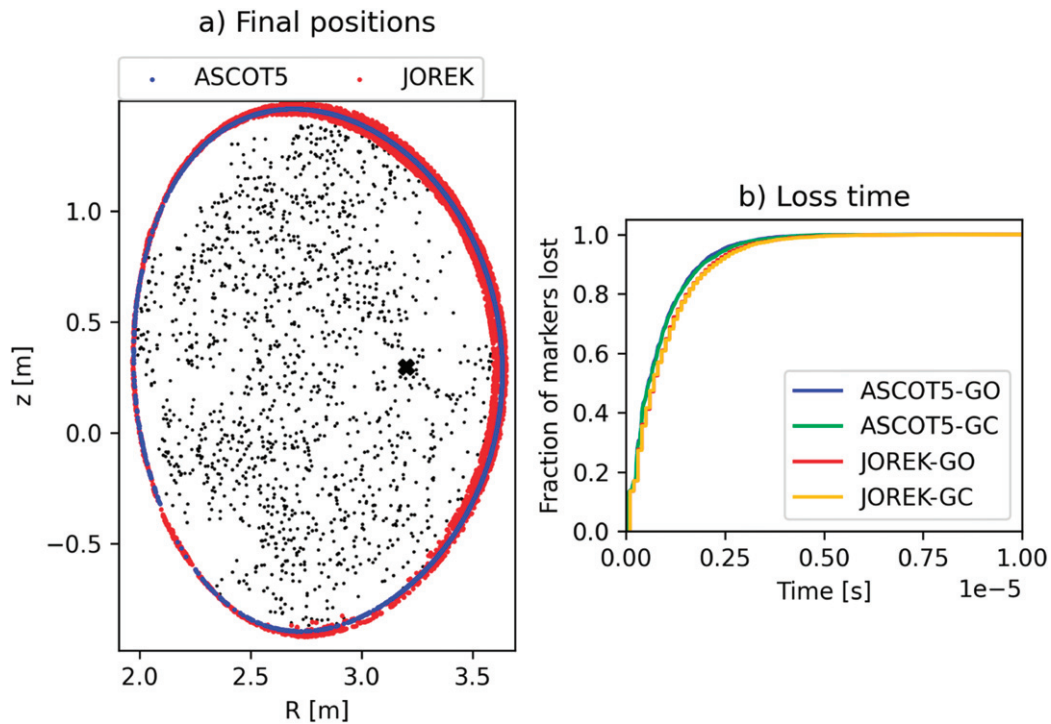
Figure 32. (*a*) The final marker position (blue and red dots) on the edge shown together with the magnetic field Poincaré plot (small black dots) and the position where markers where initialized (the black cross). (*b*) The cumulative losses as a function of time for each case. GO and GC refer to full gyro-orbit and guiding centre simulations, respectively.

by sampling the initial toroidal coordinate from a uniform distribution. The markers were traced until they exited the plasma, which was defined to occur when a marker passed a the poloidal flux value corresponding to a flux surface inside the separatrix. The accumulation of the losses in time was used as measure for transport to compare the results. For guiding centre simulations, ASCOT5 was used to calculate the guiding centre positions from the particle coordinates, and these guiding centre positions were exported to JOREK to ensure that the initial transformation did not bias the results.

The results of the benchmark are shown in figure 32. The stepwise structure seen in JOREK results is due to the simple method which had been used here in JOREK storing the marker status only at fixed time intervals, whereas ASCOT5 stores the exact loss-time of each marker at the end of the simulation. In this test case, there is little difference between the full gyro-orbit and guiding centre results in both codes. Between JOREK and ASCOT5 there is a slight deviation, and it is not certain what is the cause. One possibility is the fact that the poloidal flux in ASCOT5 is an axisymmetric quantity (the 3D components of the field are evaluated separately), and so the markers are not terminated at the same (R, Z)-coordinates as can be seen in figure 32(*b*). Nevertheless, the results show good agreement between the different pushers in both codes.

5. Applications to ELMs and ELM control

Unmitigated type-I ELMs [28, 178] cannot be tolerated in ITER at least in full current operation [179] and even ELM types associated with reduced energy losses from the plasma and lower transient heat loads to plasma facing components,

are likely not acceptable for DEMO [180]. Consequently, the dynamics of various ELM types, the access to ELM free regimes, and a robust understanding of ELM control methods are crucial research topics regarding the successful design and operation of future fusion devices. In this section, JOREK simulations regarding the plasma pedestal, edge and scrape off layer are summarized focussing on recent results. Section 5.1 presents work on natural ELM crashes and cycles, section 5.2 contains results regarding the triggering of ELMs by pellets, section 5.3 shows the excitation of an ELM crash by a vertical magnetic kick, section 5.4 addresses the control of ELMs by RMP fields, section 5.5 presents work on the simulation of ELM free regimes, and section 5.6 shows recent results on an advanced modelling of the SOL and divertor region as well as detachment physics. Section 5.7 finally, contains a very brief outlook. For investigations of ELMs and ELM control with other non-linear MHD codes, refer to the review provided in reference [3] and for more recent work to references [181–191] and references therein.

5.1. Natural ELMs

The phenomenology of ELMs has been studied intensively since the discovery of H-modes in tokamaks [192]. Several experimental observations like magnetic field, density and temperature perturbations localised on the low field side (LFS) of the tokamak suggest that MHD instabilities are involved in ELMs phenomenology [28, 178]. In particular, linear ideal MHD theory and codes indicate that ballooning modes driven by large edge pressure gradient and kink-peeling modes (KPMs) driven by large edge bootstrap current, both
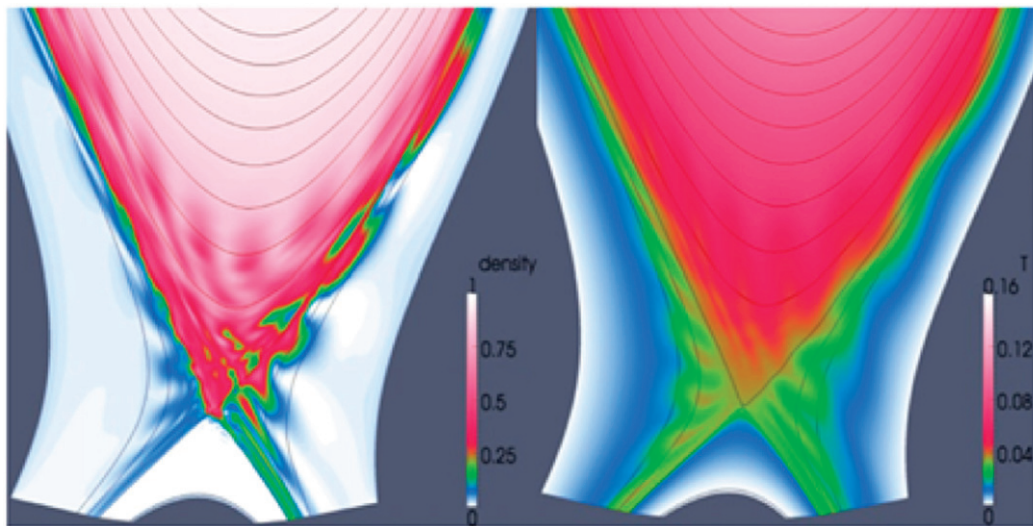
Figure 33. Perturbed density and temperature in the poloidal plane at the time of the maximum of ELM MP for a JET-like 3 MA plasma. Reproduced courtesy of IAEA. Figure from [194]. Copyright 2013 IAEA.
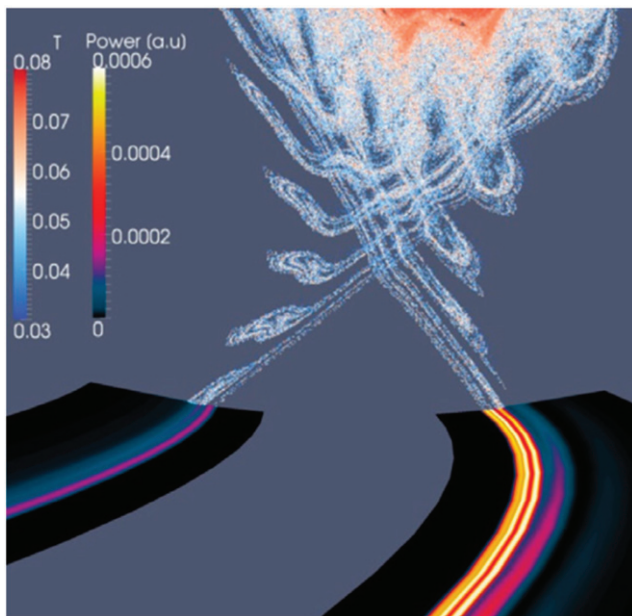
Figure 34. The homoclinic tangles forming during an ELM crash in JET and the resulting strike line splitting are shown. Reproduced courtesy of IAEA. Figure from [194]. Copyright 2013 IAEA.

typical for the pedestal region in H-mode scenarios, are the underlying physical instabilities of ELMs destabilisation [28, 29, 178]. However, only non-linear resistive MHD codes could explain the full non-linear dynamics of ELMs, reproducing in particular rotation of the modes, the mechanisms of density and temperature profiles relaxation, ELM cycling behaviour, heat and particle fluxes the heat deposition structure arriving in the divertor and many other key experimental observations during ELMs.

The first non-linear simulations of ELMs with JOREK were done without two fluid diamagnetic effects and are described in references [1, 8, 36, 149, 193] for JET-scale plasmas. Indeed, it was confirmed that ELM offset is due to the destabilisation

of coupled kink-peeling and medium-*n* ballooning modes and their linear phase is very close to the ideal MHD predictions. However, compared to the ideal linear MHD, the presence of the separatrix and plasma resistivity showed a strong stabilization of the ideal MHD external KPMs replacing them by the so-called peeling-TM which is much less sensitive to edge safety factor *q* value compared to ideal MHD description. The different transport channels for thermal energy and particles during an ELM crash and resulting heat and particle fluxes in divertor were described in reference [1, 8, 193]. The non-linear evolution of a medium-*n* ballooning mode shows the formation of density filaments and, due to non-zero resistivity, causes magnetic reconnection leading to stochastic magnetic fields. Thereafter, the density evolution is mostly determined by the $E \times B$ convection cells (figure 33, left), while the temperature evolution (figure 33, right) is mainly dominated by parallel conduction in the stochastic fields formed during the ELM crash (figure 34).

The density filaments are sheared off from the main plasma by a mean $n = 0$ poloidal flow which is non-linearly induced via Maxwell stress [8] forming 'blobs', expelled from the main plasma towards the SOL. The amplitude of the ballooning mode is limited by this mean flow stabilising it and multiple (in time) density filaments can develop to bring the plasma below the stability boundary. The importance of poloidal flows onto ELM crashes on one hand and the generation of poloidal flows during the crash on the other hand was investigated also in references [73, 149]. It was shown that large ELMs are mostly conductive-type, meaning relatively large losses in pedestal temperature, because the associated MPs cause ergodisation of the edge and lead to significant conductive energy losses along perturbed field lines. The interception of the resulting homoclinic tangles with divertor plates (figure 34) determine the pattern of the ELM heat flux in the divertor. Therefore explaining the splitting of the strike lines to the divertor during an ELM often observed in the experiments [194]. The smaller ELMs are mostly convective-type with dominant

density losses. These ELMs have divertor footprints determined by the radial distance travelled by plasma filaments expelled by an ELM and the loss of the plasma energy is determined by the energy stored in the expelled filaments. The simulated divertor wetted area during the ELM was shown to grow linearly with the ELM size in agreement with experimental observations. ITER simulations [193] showed that for ELM losses up to 4 MJ of energy, the broadening of the wetted area is similar for conductive and convective ELMs, in spite of the different origin of the widening.

After the first confirmation of the generic features of full non-linear ELM dynamics using the JOREK code, extensive studies of ELMs physics in existing tokamaks (JET, ASDEX Upgrade, KSTAR, DIII-D, MAST) were performed with the aim of further validation and improvement of JOREK physical models for more confident predictive modelling for next-step machines and in particular ITER. Such extensive studies have also enabled the JOREK code to contribute to general ELM physics issues like the role of filamentary structures in ELM and ELM control [195]. In the following, we briefly describe some examples and main results. First simulations of type-I ELMs in the ASDEX Upgrade tokamak [196] were shown in reference [197], where a spatial and temporal substructure of the ELM were observed leading to several bursts. This work pointed out the importance of including multiple toroidal harmonics in non-linear modelling to obtain a realistic dynamics of ELMs. After the early linear growth phase dominated by the most unstable medium-*n* ballooning mode, many other harmonics become destabilized through non-linear coupling while approaching the ELM crash. The non-linear destabilization of the low-*n* harmonics explained the poloidally and toroidally localized structures of the ELM crash [198] observed in experiment [199] as 'solitary MPs' (figure 35). The excitation of low-*n* harmonics during an ELM crash by non-linear mode coupling (figure 36) was studied in references [200, 201]. Note two limitations of this study: $E \times B$ and diamagnetic background flows were not included and the toroidal resolution does not include higher harmonics of the linearly most unstable mode due to computational limitations. More recent work like [202] has overcome both limitations. Density filaments and their propagation into the SOL were compared qualitatively to the experiment in reference [203]. More recent simulations for ASDEX Upgrade were performed at realistic experimental parameters and with plasma background flows consistently included [5]. Here, ELM losses, magnetic measurements [204], and measurements for cold-front penetration [205] agree well with experiments. Based on such simulations, the role of density fluctuations onto ECE-imaging measurements was studied [206] and the excitation of parametric decay instabilities during electron cyclotron resonance heating was investigated [207]. Further simulations confirmed the trend of decreasing dominant toroidal mode numbers with an increasing edge safety factor like experimentally observed [208].

The two fluid diamagnetic (for electron and ions), neoclassical and toroidal background plasma flows were implemented and applied to ELM simulations in reference [209] entailing an important improvement of the physical model. A stabilizing
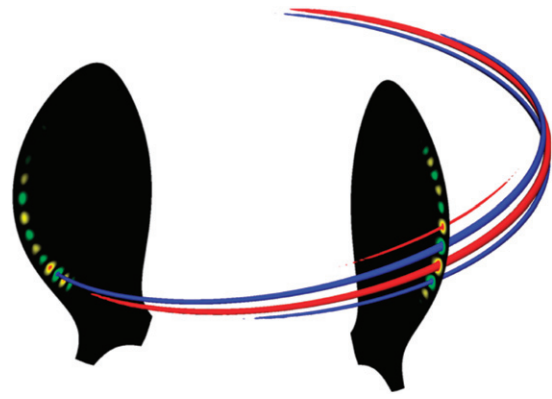

Figure 35. The poloidally and toroidally localized magnetic flux perturbation in an ASDEX Upgrade simulation is shown, which has similarities to so-called solitary MPs observed experimentally [199]. Reprinted from [199], with the permission of AIP Publishing.
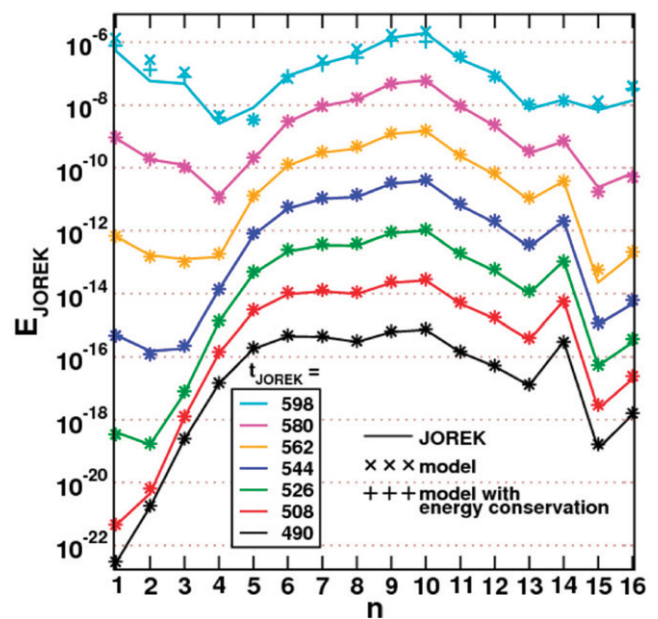
Figure 36. For various time points during a JOREK simulation (times in JOREK units), the toroidal mode spectrum is shown. The linearly most unstable mode numbers around $n = 10$ give rise to a broad mode spectrum in the non-linear phase at the ELM onset driven by non-linear mode coupling. In particular the linearly stable $n = 1$ mode is driven to amplitudes comparable to the linearly dominant $n = 10$ mode. Reprinted from [201], with the permission of AIP Publishing.

effect of poloidal flows and in particular diamagnetic effects known from the ideal MHD model onto high-*n* ballooning modes was confirmed. Moreover, after the first ELM crash, poloidal flows tend to damp the otherwise continuing ballooning mode turbulence, so the pedestal profiles can be rebuilt again by heating and particle sources until they again reach the stability limit producing the following ELM crash. The numerically obtained multi-cycle ELMs exhibit similarities to high frequency ELMs in experiments (figure 37). Moreover, the diamagnetic drifts were found to yield a near-symmetric ELM power deposition on the inner and outer divertor target plates, consistent with experimental measurements [210].
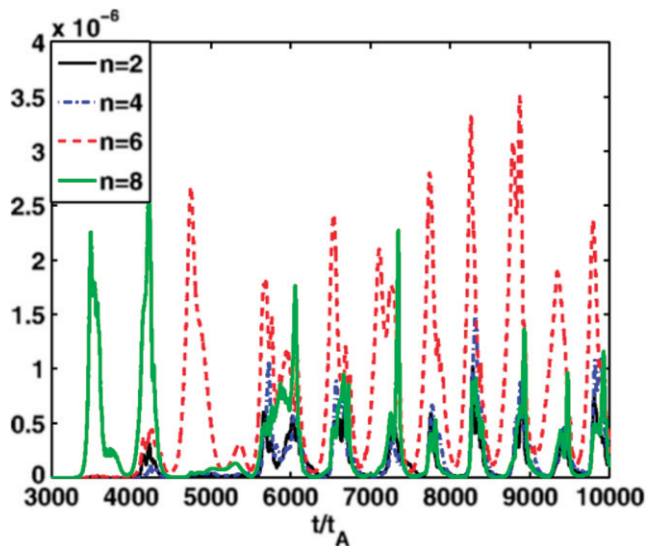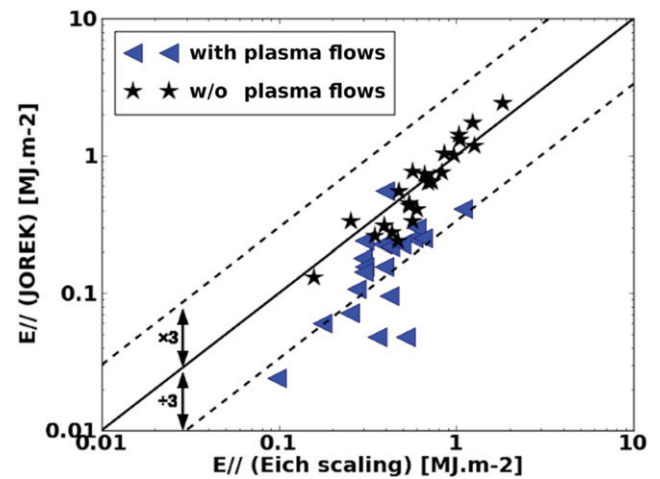
Figure 37. High frequency multi-ELM cycles were obtained with JOREK when background plasma flows were consistently included. Reprinted figure with permission from [210], Copyright 2015 by the American Physical Society.


Figure 38. The parallel heat fluence to the divertor targets is plotted versus the value predicted by the Eich scaling. Good agreement is obtained in general, while the heat fluence is underestimated in tendency, when plasma background flows are taken into account (blue triangles). Note a misprint in the original publication, which was corrected in the figure shown here: the labels for the black and blue points had accidentally been swapped both in figures 4(*a*) and (*b*) of reference [45]. Reproduced courtesy of IAEA. Figure adapted from [45]. © EURATOM 2017.

The ballooning mode rotation obtained in JOREK simulations has been successfully validated against analytical linear computations in reference [160]. In successive studies, the explanation of the rotating structures in the pedestal region before type-I ELM crashes and in the inter-ELM periods (ELM precursors) observed in the KSTAR tokamak [211] was proposed [161, 212]. The two fluid diamagnetic effects and toroidal rotations included in the model were found to be the most important factors in explaining the experimentally observed rotating structures [161].

Simulations aiming to capture ELMs in the JET tokamak [213] as realistically as possible were shown in reference [214]. The dynamics of filaments, divertor heat fluxes and the influence of collisionality onto the ELM size was studied and precursor modes were shown in ELM simulations. Successively, reference [215] addressed simulations for a series of JET ITER-like wall discharges to perform quantitative comparisons with the experiments. It was shown that the accuracy of the pre-ELM equilibria used for the simulations as well as realistic parallel heat conductivity are both critical for reproducing experimentally observed ELM energy losses. Divertor peak heat fluxes were still underestimated and the ELM duration overestimated in the simulations. Reference [45] took more accurate pre-ELM equilibrium reconstructions as basis for the simulations and also investigated the impact of $\mathbf{E} \times \mathbf{B}$ and diamagnetic background flows onto the ELM dynamics. Very good agreement with the experimental scaling law [194] for the energy fluence to the divertor targets was obtained in general (figure 38), although the values were slightly underestimated when $\mathbf{E} \times \mathbf{B}$ and diamagnetic background flows were taken into account, in particular for configurations corresponding to small ELM sizes in the experiment. Agreement regarding ELM energy losses and divertor peak heat fluxes improved significantly compared to reference [215], also here with some deviation when stabilizing background flows are taken into account. Resolving the remaining

deviations might require simulations of complete ELM cycles. Reference [45] also demonstrated, that the threshold for the onset of edge instabilities in non-linear JOREK simulations is in better agreement with experimental observations than predictions from linear codes, which typically neglect many non-ideal effects.

Using the JOREK code, realistic simulations of multiple type-I ELM cycles were recently obtained for the first time [202] (figure 39). The inclusion of pressure-gradient driven diamagnetic drifts is imperative to obtain the cyclical dynamics. The simulated ASDEX Upgrade plasma corresponds to a moderate triangularity and high pedestal density. Simulations for different plasma parameters are the subject of ongoing efforts. The simulations find that with increasing heating power, the ELM frequency rises, consistent to experimental observations of type-I ELMs [178]. A precursor-like mode activity is observed before the violent onset of the ELM crash. It was found that the interaction between this precursor-like mode and the background plasma is responsible for the explosive ELM onset. The first ELM crash simulated was observed to be different from the subsequent ELMs because of differences in the non-axisymmetric seed perturbations that exist before each ELM crash. For the first ELM, these are simply noise-level perturbations, while for the subsequent ELMs the seed perturbations are consistent with the prior existence of an ELM. This observation stresses the importance of simulating the entire ELM cycle—particularly for predictive simulations. These simulations were also used as basis for investigations regarding the scattering of ion cyclotron waves by filaments and ELMs [216] and for recent studies of pellet ELM triggering (see section 5.2).

The regime of small ELMs is another topic of current research with the JOREK code. Experiments with small
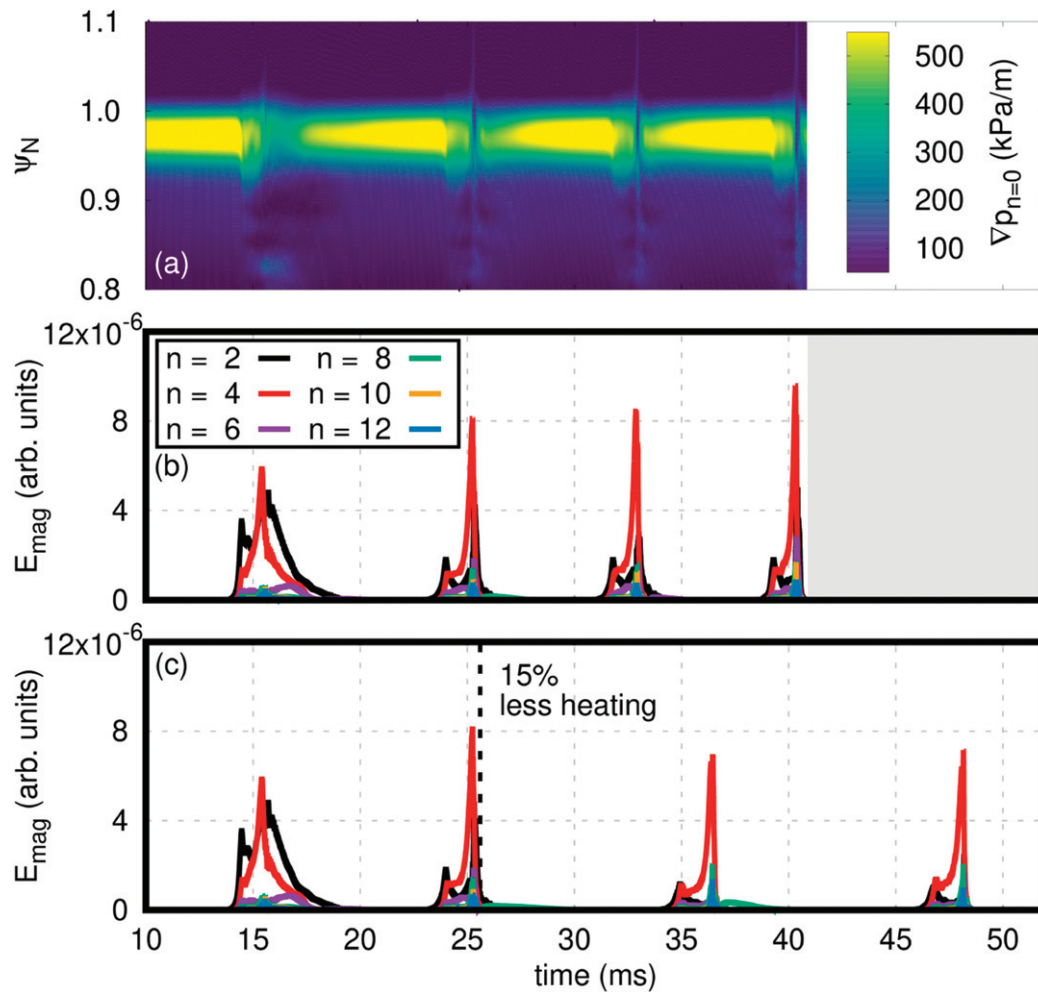
Figure 39. Results from a simulation of several type-I ELM cycles in ASDEX Upgrade adapted from reference [202]. (*a*) shows the evolution of the toroidally averaged outer midplane pressure gradient versus time and normalized poloidal flux. (*b*) contains the evolution of the magnetic energies of the $n = 2, 4, \ldots, 12$ toroidal harmonics included in the simulation. Finally, (*c*) shows the magnetic energies of a simulation with 15% less input heating power starting from 25.2 ms. The ELM repetition frequency is seen to have a direct relation with the input heating power. The first ELM crash is determined by random seed perturbations and therefore shows different dynamics from the following ones, which are fully self-consistent. Strong precursor-like modes which already affect the pedestal gradient are observed prior to the very fast crashes.

ELMs host a quasi-continuous power exhaust and, with appropriate plasma shaping, can avoid type-I ELMs completely [217]. Using the simulation set-up for the type-I ELM cycles described above, but with a lower heating power, an operational regime with small ELM-like behaviour has been also obtained [218], albeit not yet at ITER relevant conditions which remain a topic for future research. These simulations highlight the importance of the stabilising effects of plasma flow onto small ELMs [219].

Based on ELM simulations, fast ion losses were studied by particle tracing in the fields of the JOREK simulations confirming increased losses during the ELM crash [220–222]. Tungsten transport during a large ELM crash in ASDEX Upgrade was investigated kinetically [223] (figure 40) revealing that lower-dimensional diffusion–convection models cannot explain the far reaching interchange driven transport of the impurities and that high tungsten densities in the ITER SOL might cause a net inward transport of tungsten during an ELM crash.

Simulations for double X-point plasmas in the spherical tokamak MAST [224] were reported in reference [225]. In particular, the filament dynamics were compared to experimental observations revealing good qualitative agreement regarding structure and dynamics (figure 41). The general role of filaments in ELM dynamics and ELM energy losses was also addressed in reference [195]. The ELM energy losses, and divertor heat flux profiles were found to agree reasonably well in spite of using a reduced MHD model in this low aspect-ratio configuration and neglecting diamagnetic drift effects. For the MAST Upgrade tokamak, predictive simulations were reported in references [76, 159, 163, 226]. Particular attention was paid to studying the effect of the Super-X divertor onto detachment in the inter-ELM phase, and burn-through during the ELM crash and resulting divertor heat fluxes. More
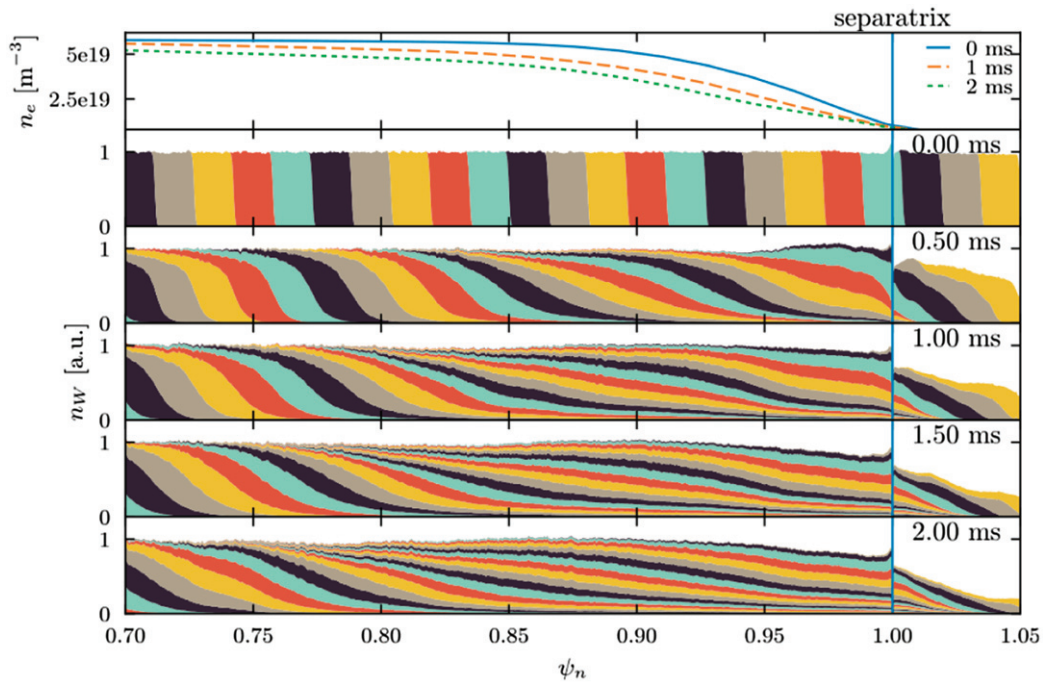
Figure 40. During an ELM crash in ASDEX Upgrade (see reference [5]), the transport of tungsten is investigated. The top figure shows the evolution of the density profile during the ELM crash, the figures below show the mixing of tungsten particles due to the MHD activity. Particles are coloured according to their original location. A strong mixing is observed including a strong interchange transport across the separatrix. Reprinted from [224], with the permission of AIP Publishing.
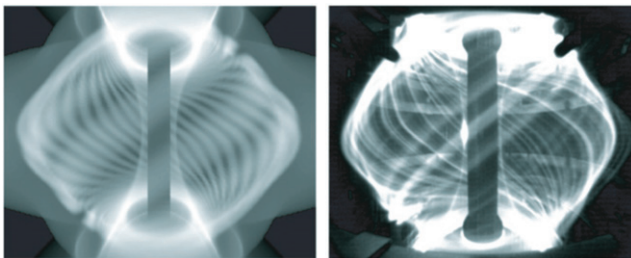
Figure 41. The signature of ELM filaments in visible light is shown for ELMs in MAST. Left: virtual diagnostic data based on a JOREK simulation. Right: visible light picture in the experiment. Reproduced from [226]. © IOP Publishing Ltd. All rights reserved.

details on this study are given in section 5.6 along with detachment and burn-through studies for ITER reported in reference [77].

5.2. Pellet ELM pacing

Pellet injection into the pedestal of H-mode plasmas has experimentally shown to be capable of increasing the ELM frequency significantly above its natural value while decreasing ELM sizes, which may be beneficial for divertor lifetime and, in particular, for the control of impurity concentrations in the plasma. In non-linear simulations, such ELM triggering has first been demonstrated in references [8, 227]. Detailed comparisons to experimental observations for pellet triggered ELMs in the DIII-D tokamak were reported in reference [86]. These studies revealed the detailed physical mechanisms for the destabilization of the ELM by a pellet: ablation leads to a largely adiabatic increase of the plasma density and decrease of the plasma temperature in the vicinity of the pellet. Due to the high mobility of the electrons along field lines, the temperature within a flux surface is equilibrated again on an extremely fast time scale, while the density transport takes place on the far slower time scale of the ion sound speed. In combination, this leads to a helical structure with strongly increased pressure around the pellet location, which locally exceeds the ballooning threshold. In the further evolution, the initially localized perturbation spreads poloidally and toroidally such that the actual ELM crash leads to radial transport across the separatrix in a far less localized manner. A clear threshold was observed between a perturbation of the plasma and ELM triggering (figure 42).

Simulations for ELM triggering in the JET tokamak revealed a toroidal asymmetry in the peak heat flux to the divertor [228, 229], which constitutes a clear difference between the non-linear behaviour of a natural and pellet triggered ELM. The simulation shows good quantitative agreement (within 10%–20%) against the divertor heat flux obtained with experimental thermography measurements at the single toroidal location where measurements are available. However, simulations predict a significantly larger divertor heat flux at a location which was not covered by measurements in the corresponding experiments.

Recent simulations for pellet triggered ELMs in ASDEX Upgrade [230] are studying the ELM triggering possibility at different phases of the pedestal build up, which provides additional insights into the evolution of the pedestal stability in the inter-ELM phase. These pellet simulations include self-consistent $E \times B$ and diamagnetic background plasma flows,
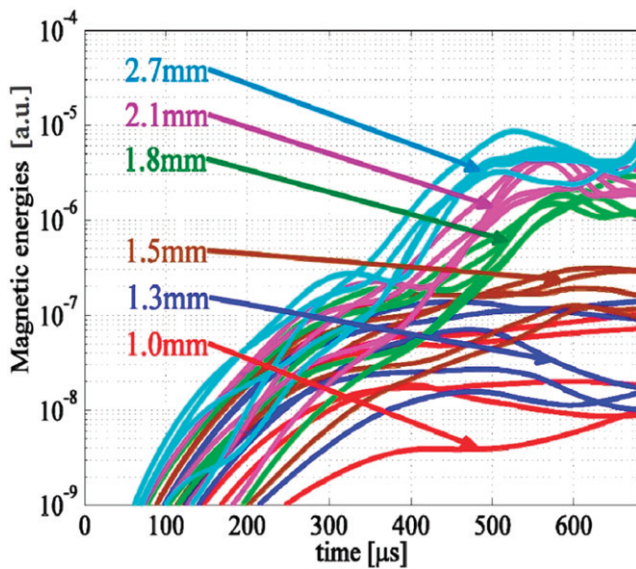
Figure 42. The excitation of MPs during pellet injection is shown. At a pellet size between 1.5 mm and 1.8 mm, a bifurcation is observed. For larger pellet sizes, a strong and sudden excitation of MP energies to higher amplitudes is observed, which is absent for the smaller pellet sizes. These strong perturbations are associated to a pellet triggered ELM crash and induce losses from the plasma. Reproduced courtesy of IAEA. Figure from [87]. Copyright 2014 IAEA.

use realistic plasma parameters and are based on the type-I ELM cycle simulations of reference [202]. The injection of pellets during the pedestal build-up is studied varying injection time, pellet size, and pellet velocity. Figure 43 shows the divertor heat flux versus time and the divertor length (at the toroidal position of the pellet injection) for three cases: pellets with 0.8×10^{20} D atoms are injected into the plasma at 8, 10, and 12 ms during the pedestal build-up with an injection velocity of 560 m s^{-1}, while the natural ELM crash would occur at around 16 ms. Note that the number of atoms reaching the plasma is given here; thus the simulation corresponds to a pellet with approximately 1.5×10^{20} atoms in the experiment (assuming 50% losses in the guide tube). A sharp transition is observed between 10 and 12 ms from no-ELM response to ELM triggering. Thus, the experimentally observed [231] lag-time has been reproduced qualitatively here, during which no ELM triggering is possible by means of pellet injection. In case of ELM triggering, large transient heat fluxes are observed in the range of ~ 20 MW m^{-2} for about 0.4 ms. Furthermore, the ELM triggering cases show a far broader mode spectrum compared to the no-ELM response and the triggered ELMs features a toroidally asymmetric heat deposition with a strong $n = 1$ component. A detailed comparison of the non-linear properties of triggered and spontaneous ELMs, e.g., regarding the wetted area, are shown in reference [232].

5.3. ELM control by magnetic kicks

Besides pellet injection discussed in the previous section, ELM pacing has experimentally also been demonstrated via so-called 'vertical magnetic kicks', during which the current in

one or several poloidal field coils is evolved in time in such a way, that the plasma undergoes an excursion above and below the original location.

A benchmark of the plasma excursion caused by a variation of the poloidal field coils (figure 27), and fully self-consistent simulations of ELM triggering by vertical magnetic kicks were shown in references [92, 171, 233] for an ITER 7.5 MA plasma (figure 44). Consistently with the observations from several fusion experiments, it was demonstrated that PBMs are triggered during a downward (towards the X-point) excursion of the plasma leading to an ELM crash in an otherwise stable plasma configuration, while an upward excursion of the plasma does not give rise to such edge instabilities. The edge instabilities were shown to always appear at a particular vertical displacement of the plasma independently of the actual time scale of the applied oscillation, like also seen experimentally. Detailed analysis of the simulations and comparisons to analytical considerations allowed to confirm that an increase of the plasma edge current density during the downward motion of the plasma in the inhomogeneous magnetic field is responsible for the destabilization of the PBMs. A detailed analytical picture of the mechanisms of the edge current evolution was obtained consistent with the simulation results.

5.4. ELM control by RMPs

Simulations for the penetration of external MPs into plasmas of the DIII-D tokamak were shown in references [37, 234] for the first time, demonstrating that the magnetic topology obtained in simulations with plasma rotation considerably differs from the so-called 'vacuum approximation', in which the vacuum magnetic field of the perturbation coils is simply added to the equilibrium magnetic field of the plasma (figure 45). A strong suppression of magnetic islands and stochastic field regions by the plasma was shown in these simulations carried out in the 'zero-beta' limit and with a rigid-body rotation. A radial $E \times B$ convective transport was observed in the presence of the MP fields and an important role of it for the experimentally observed density pump-out was proposed. However, note that the density pump-out due to this mechanism was not strong enough compared to the experiment.

With the implementation of neoclassical, diamagnetic and toroidal background plasma flows and taking into account also the self-consistent evolution of the plasma temperature, references [49, 235, 236] were able to show in detail the dynamics of the penetration of the RMPs into the plasma. In particular, the screening effect by the rotating plasma was studied in detail. Based on such simulations, reference [237] investigated the three dimensional lobe-structure of the homoclinic tangles formed in the presence of the 3D magnetic field perturbations in detail revealing a strike-line splitting like it is also observed in many experiments [238]. References [235, 236] also investigated the effect of the RMP fields onto the edge instabilities. In particular, it was shown that strong RMP fields can suppress peeling-ballooning instabilities in the simulations replacing them by saturated modes (figure 46), and that
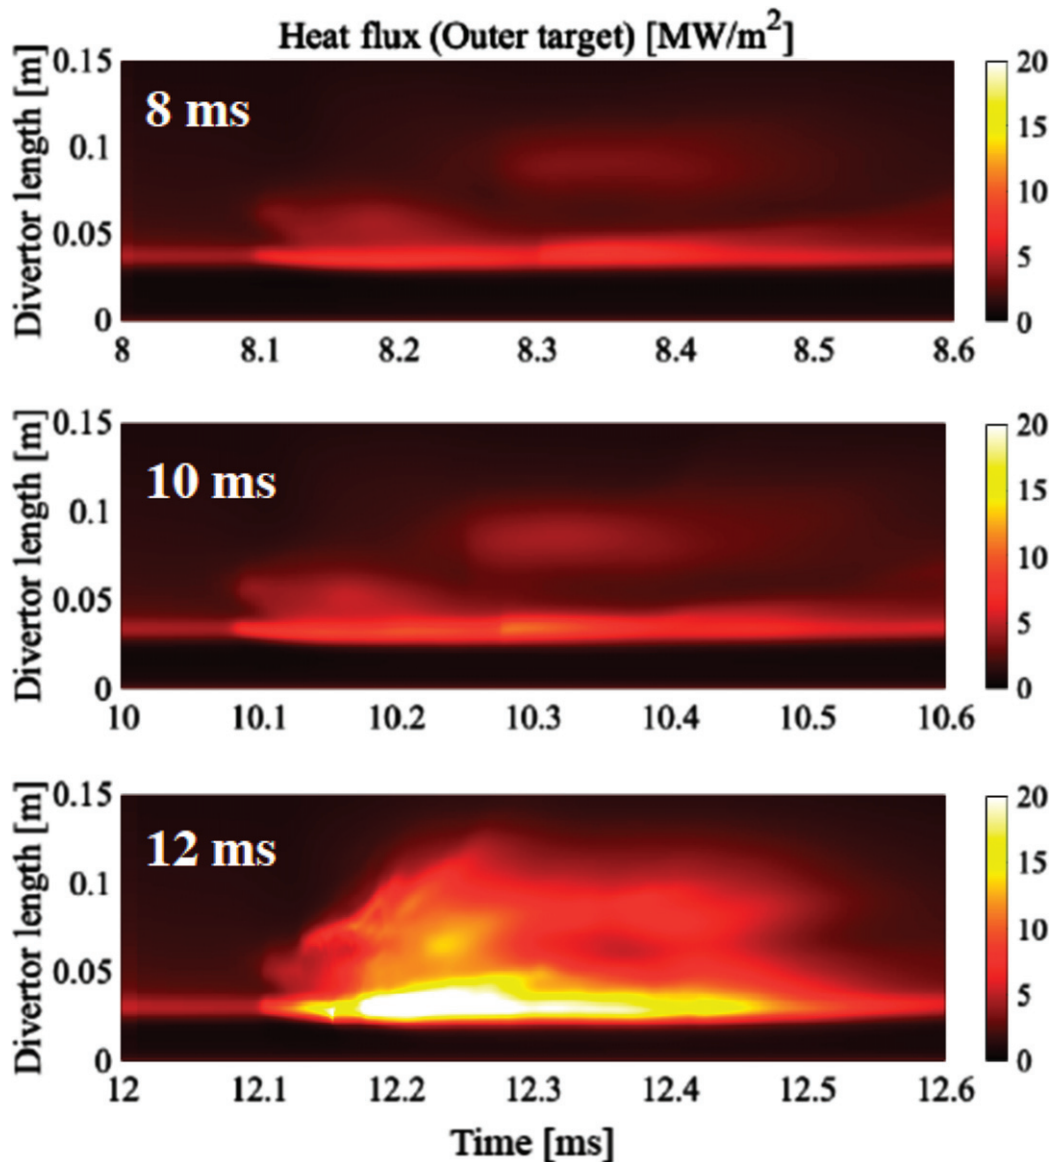
Figure 43. The time evolution of the heat flux onto the outer divertor targets which is caused by 0.8×10^{20} D pellet injection with 560 m s^{-1} is shown. Three cases are compared, where the transition from no-ELM response (8 and 10 ms) to ELM triggering (12 ms) is observed. Reproduced courtesy of IAEA. Figure from [231]. © EURATOM 2021.

uneven-*n* modes remain strongly sub-dominant in the presence of an $n = 2$ MP, even if these are linearly unstable in the absence of the perturbation fields.

Furthermore, simulations for RMP experiments in ASDEX Upgrade showed very good agreement for the penetration of the external fields into the plasma and the resulting corrugation of the separatrix for several investigated 'coil current phases', i.e., different perturbation spectra [239]. The configuration leading to the strongest ELM mitigation effect in the experiments was identified in the simulation as the one with the largest kink response of the plasma near the X-point. In further studies shown in references [240, 241], the interaction of the RMP fields with plasma edge instabilities was investigated in a plasma configuration leading to an ELM crash in simulations without RMP fields. When increasing

the RMP amplitude at given plasma rotation, and also when reducing the plasma rotation at fixed RMP amplitude, a transition was observed from an unmitigated ELM regime into a mitigated ELM regime with reduced perturbation amplitudes and losses, and further into an ELM suppressed state. While the 3D perturbations observed in the unmitigated and mitigated state show a 'bursty' behaviour and rotate in the electron diamagnetic direction, the suppressed state is characterized by saturated modes, which are not rotating in the lab frame due to a locking with the external perturbation fields (figure 47). Non-linear mode coupling is shown to be crucial for the ELM suppression.

Recent simulations for KSTAR [212, 242] investigate a plasma configuration unstable with respect to PBMs. When $n = 2$ RMPs are applied, a formation of islands and a stochas-
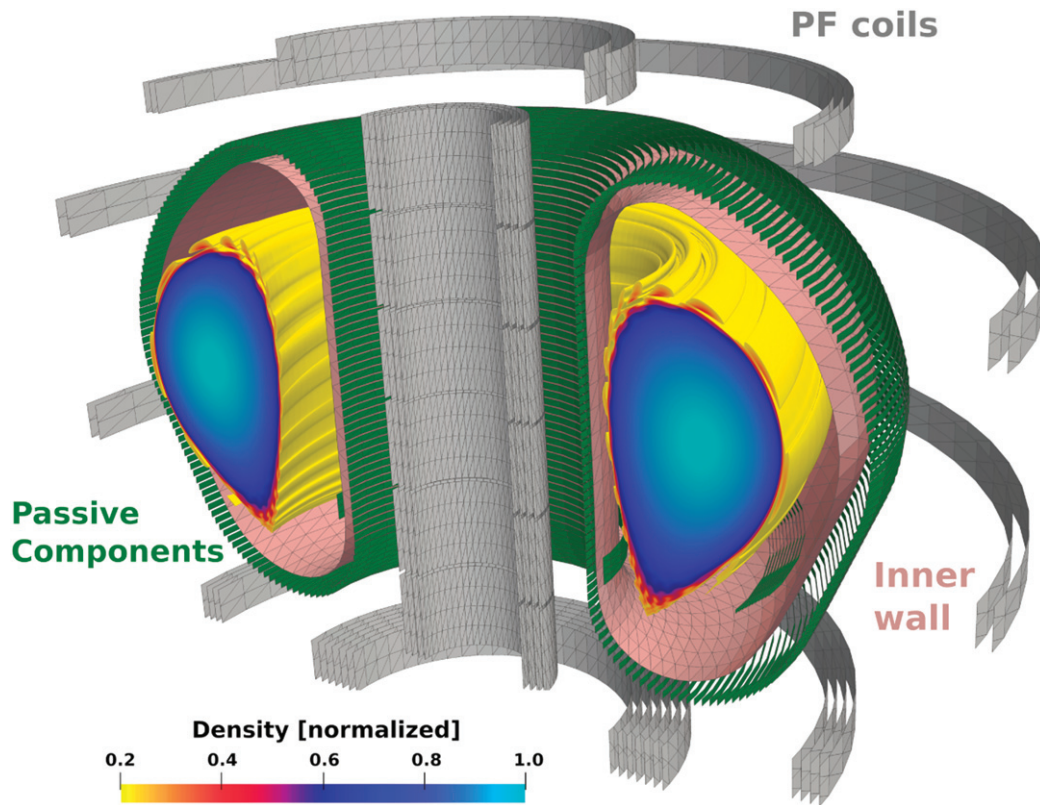
Figure 44. The perturbed density during an ELM crash induced by a vertical magnetic kick is shown. All relevant conducting structures were included in the JOREK–STARWALL simulation. The figure was generated based on the simulation data published in reference [171].
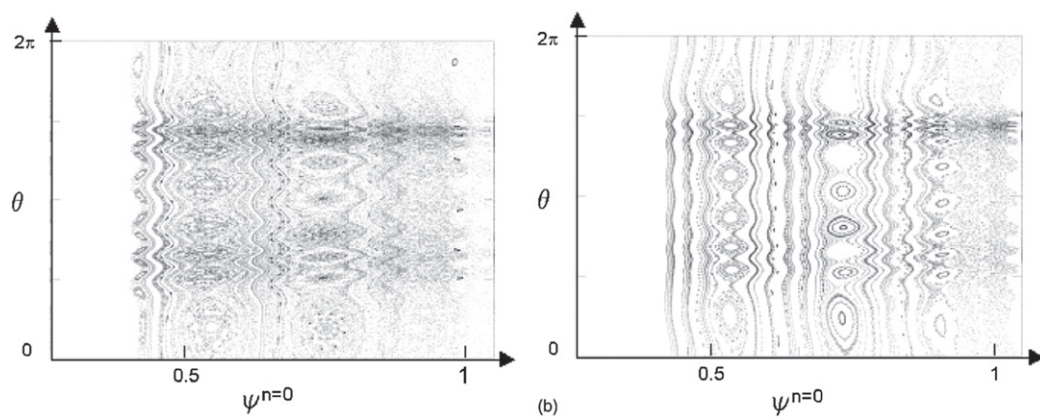
Figure 45. The magnetic topology caused by RMPs is shown. In the left figure, representing the 'vacuum field' perturbation many island chains and stochastic regions are visible. In the right figure which includes the plasma response, most resonant components are shielded by the rotating plasma, and only a few penetrated island chains are visible. Reprinted from [37], with the permission of AIP Publishing.

tic field layer is observed at the plasma boundary and a density pump-out is observed, while the effect onto the temperature is weaker. The article shows, that the PBMs can be mitigated and eventually suppressed by the application of the RMP fields of sufficient amplitude. Consistently with reference [240], a breaking of the mode rotation is seen in the suppressed state. Non-linear mode coupling plays an important role in the suppression, since the PBMs would be still unstable in the degraded plasma pedestal, when the 3D fields are removed from the simulation.

5.5. ELM free regimes

Non-linearly saturated external kink modes in X-point plasma simulations including the SOL were described for the first time in reference [36]. Dedicated simulations for quiescent H-mode (QH) plasmas from DIII-D were carried out in reference [243, 244]. These showed the development of saturated KPMs leading to the characteristic edge harmonic oscillation (EHO), which is observed experimentally in the density evolution and is caused by the rotating saturated modes
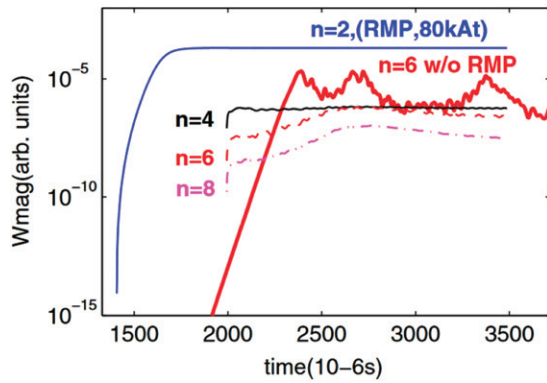
Figure 46. The MP energies of a natural ELM crash (simulated with a single toroidal mode number $n = 6$ are compared to a simulation with applied RMPs. The $n = 6$ MP energies are significantly reduced in the presence of the RMP fields. Instead of the bursts corresponding to an ELM crash, stationary modes are formed. Reprinted figure with permission from [236], Copyright 2014 by the American Physical Society.

(figure 48). The influence of shear flows and resistive wall effects onto the development of the QH-mode was studied for DIII-D and predictively for ITER. Linear stability studies for the DIII-D case showed that the case is close to the peeling stability boundary. It was demonstrated that the saturated KPMs are replaced by an ELM crash, if the edge current density is decreased and the pedestal pressure increased.

Recently, simulations for experiments in ASDEX Upgrade have also been performed [245]. A QH-mode experiment of the carbon wall ASDEX Upgrade was studied and saturated low-n KPMs were found to be non-linearly dominant at the edge of the plasma forming a helically perturbed structure. This is in spite of an initial small crash triggered by high-n ballooning modes, confirming that the simulations predict the formation of a QH-mode like state for this plasma configuration. In the same reference, also simulations for the ITER $Q = 10$ baseline scenario were performed. Although further investigations may be necessary, these simulations indicate, that the bootstrap current may be high enough in this scenario to enter a QH-mode regime. The $E \times B$ flows are shown to have a stabilizing influence on high-n modes in the ITER plasma. In the ITER relevant range of $E \times B$ rotation around 20 km s^{-1} at the pedestal, the simulations show that low-n modes $n = 1$–3 are destabilized and become dominant in the non-linear saturated state that establishes, while higher-n modes $n = 7$–10 are significantly suppressed. The saturated low-n modes lead to density oscillations in the pedestal which are typical for EHO characteristics. Also, simulations with the free boundary extension (section 2.9) were shown in order to investigate the effect of the true geometry of the vacuum vessel onto the instabilities. Related to the effect of resistive walls onto the KPM, simplified resistive wall mode studies had already been performed and benchmarked in references [168, 246] (figure 49).

5.6. Detachment physics

The simulations presented in references [76, 77] study detachment/burn-through during an ELM crash in a non-linear MHD simulation. Using the neutrals fluid model (section 2.6), simulations of plasma detachment for the MAST Upgrade tokamak were carried out [76, 163] and reasonable agreement with SOLPS simulations could be shown, although some discrepancies will require further investigations. In particular, the characteristic 'roll-over' could be qualitative reproduced, where the target particle flux first increases and then decreases for an increasing midplane separatrix density, leading to a target temperature of only a few electron volts (eV) in the detached state. Simulations of large type-I ELM crashes were carried out in such a detached plasma state, showing the 'burn-through' and re-attachment in the Super-X divertor leg. The ELM energy fluence was found to be significantly lower than that predicted from the empirical scaling for the detached Super-X divertor [163]. Similar physics was studied for ITER [77]. It could be shown that small amplitude ELMs are sufficient to re-attach the plasma in the ITER high recycling divertor transiently, increasing the electron temperature at the divertor target from a few eV to several hundred eV within a fraction of a millisecond (figure 50). In this reference, also an outlook to refined modelling of detachment and burn-through using a kinetic neutrals model (based on the particles framework described in section 2.10) is given, which will allow to capture the SOL dynamics even more accurately.

5.7. Outlook

Regarding the modelling of edge instabilities, further emphasis will be put onto a more accurate modelling of the SOL and divertor physics like shown for instance in section 5.6. Simulations of the edge plasmas including such effects, further model enhancements, and based on further numerical improvements, aim to extend the modelling of ELM cycles to fully realistic plasma parameters and describe the physics mechanisms relevant for the transitions between the regimes of different ELM types and ELM free regimes.

6. Applications to disruptions and their control

In view of ITER, disruptions are presently the highest priority topic when it comes to large-scale plasma instabilities. Unmitigated disruptions are considered intolerable above modest values (by ITER standards) of the plasma current and thermal energy. The ITER DMS, which is planned to rely on SPI, needs in particular to avoid substantial damage from heat loads, RE impacts and electromagnetic forces [247]. The current strategy to achieve this consists in: (1) mitigating heat loads by dissipating most of the plasma energy through uniform radiation, (2) avoiding, if possible, the formation of an RE beam by raising the electron density, (3) should an RE
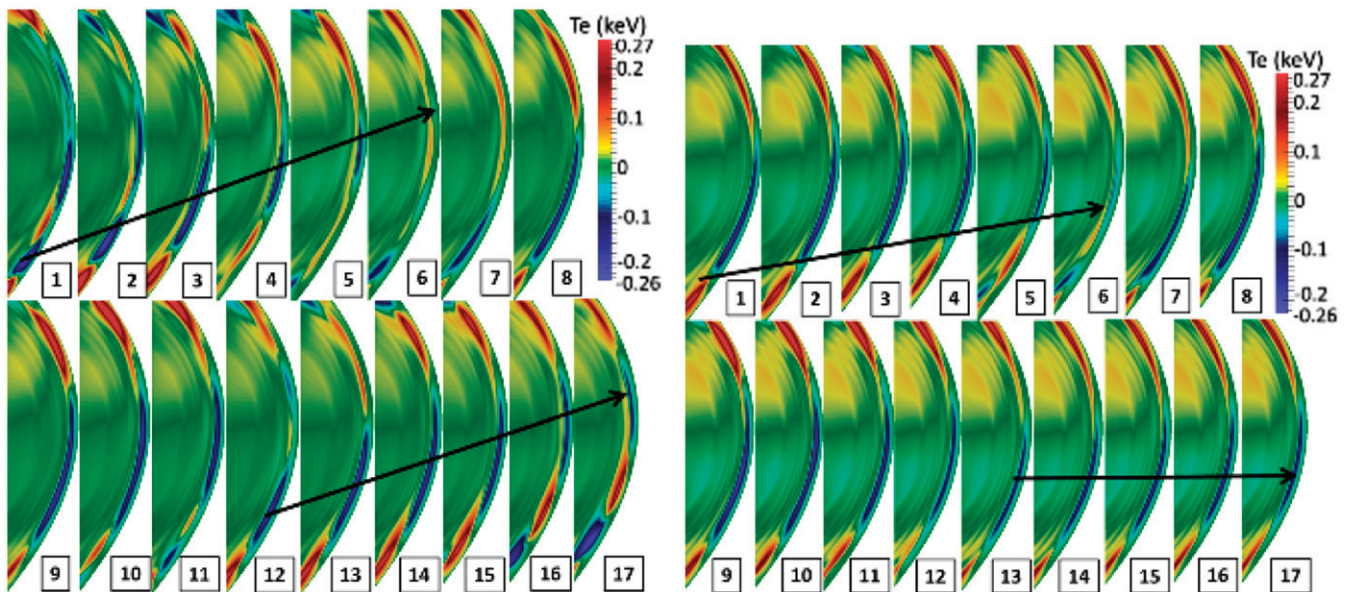
Figure 47. Edge temperature perturbations (non-axisymmetric component) for ASDEX Upgrade simulations of a peeling-ballooning unstable plasma with different RMP amplitudes. The left panel shows different time points during a simulation with moderate RMP amplitude. Edge instabilities form rotating structures of reduced amplitudes compared to the natural ELM crash (mitigation regime). The right panel shows, that a higher RMP amplitude causes a locking of the instabilities. In this case, losses from the plasma are strongly reduced and stationary mode activity is seen (suppression regime). Reprinted from [240], with the permission of AIP Publishing.
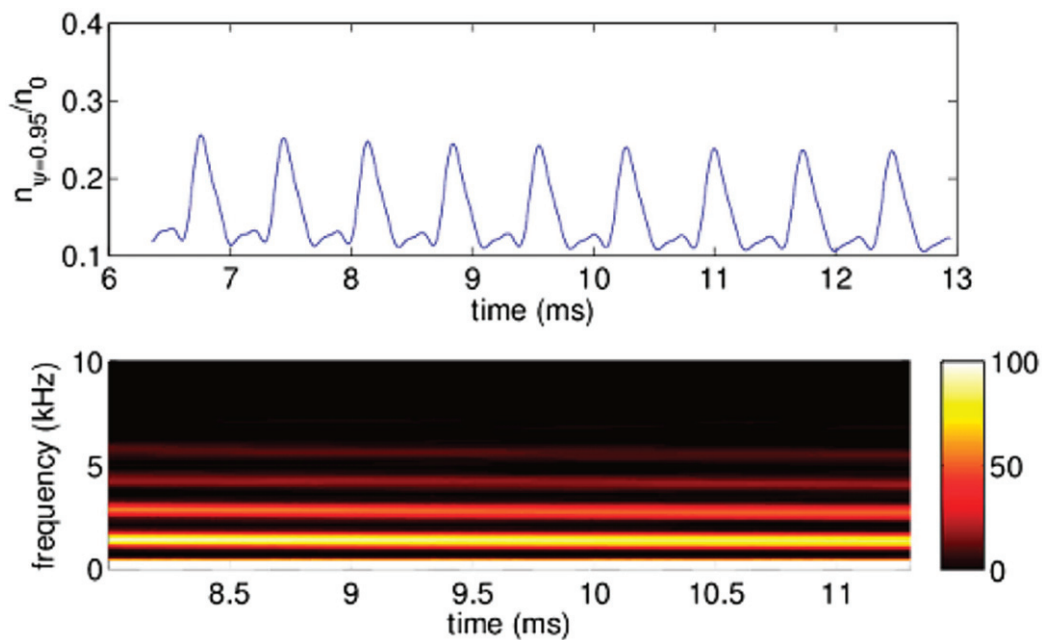
Figure 48. The edge density perturbation during a QH-mode simulation is shown. The upper figure contains the time trace of the perturbation in the outer midplane induced by a stationary and rotating edge mode, the lower figure contains the spectrum analysis of this signal. The results are in very good agreement to the experimentally observed EHOs. Reproduced from [244]. © IOP Publishing Ltd. All rights reserved.

beam appear anyway, using SPI into the beam in order to make its impact as benign as possible, and (4) controlling (via the plasma impurity content) the $I_{\rm p}$ decay rate in order to mitigate electromagnetic loads. Achieving these goals simultaneously requires a deep understanding of SPI and disruption physics, and this motivates the many disruption-related investigations

with JOREK which are described in this section. In particular, section 6.1 discusses pre-disruption physics, i.e., the mechanisms leading to disruptions, section 6.2 covers the dynamics of disruptions triggered by MMI, section 6.3 deals with VDEs and halo current dynamics, and finally section 6.4 describes RE studies.
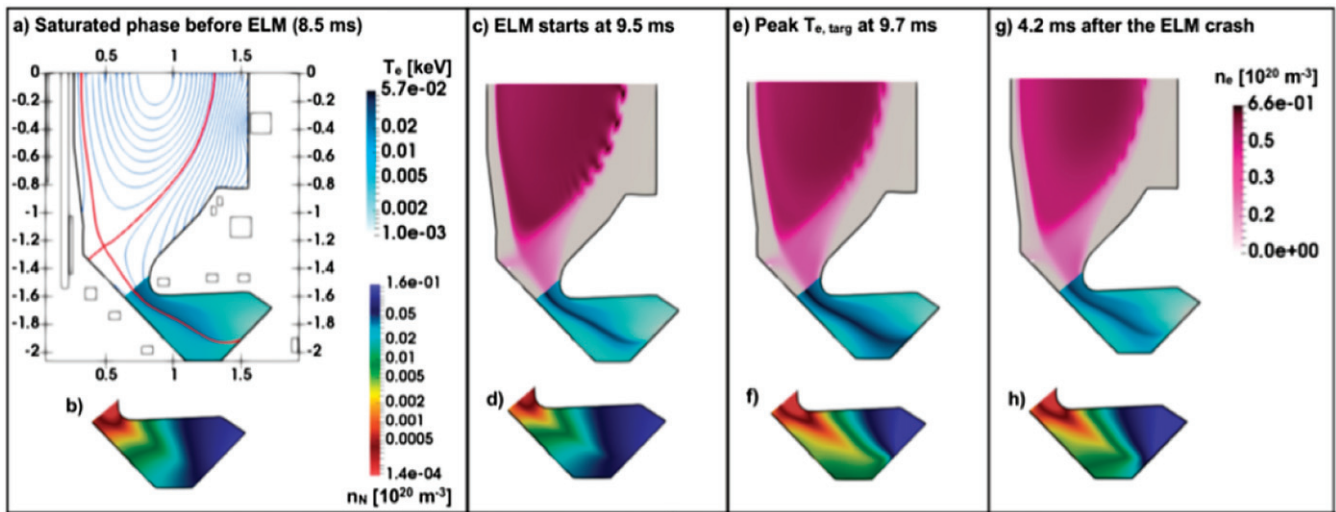
Figure 49. The simulation of an ELM crash in MAST Upgrade is shown. The simulations are performed using the neutrals model available in JOREK. The originally detached state with low target temperatures (left plot) is re-attaching during the ELM crash in the Super-X divertor (right plots). Reproduced from [77]. CC BY 3.0.
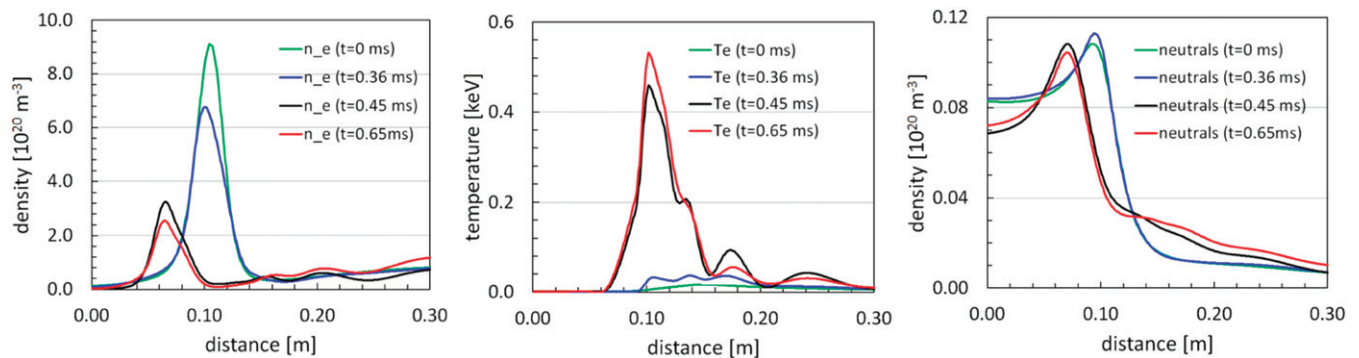
Figure 50. Density, temperature and neutral density profiles along the ITER outer divertor. Reproduced from [78]. CC BY 3.0.

6.1. Pre-disruption physics

6.1.1. Tearing mode dynamics and mode locking. Simulations in ASDEX Upgrade-like geometry have been run in which the current profile was tailored to make a dominantly $m/n = 2/1$ TM strongly unstable [248]. Due to the inhomogeneous magnetic field as well as geometrical effects, the $n = 1$ mode contains sidebands, i.e., components of the type $m/1$ with $m \neq 2$. In the linear phase, Poincaré cross-sections thus show the presence of not only a 2/1 island, but also 3/1, 4/1 and 5/1 islands. In the non-linear phase, $n > 1$ islands (in particular 3/2 and 5/2) grow fast. An analysis of the growth rate of the islands width, shown in figure 51, suggests that $n > 1$ islands result from mode coupling. For example, the growth rate of $n = 2$ islands is twice that of $n = 1$ islands. Island overlapping leads to a stochastization of the magnetic field over roughly the outer half of the plasma. Due to the fast parallel heat transport, the temperature profile flattens across this region, reproducing a key feature of so-called 'partial TQs' in experiments [249].

The amplitude of the magnetic field perturbation at the onset of the partial TQ, as measured by synthetic magnetic
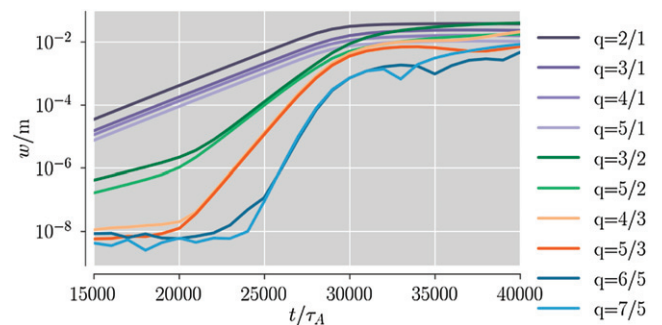
Figure 51. Island width as a function of time for various islands, showing that higher n islands are driven by mode coupling.

sensors localized in the midplane, was compared to the empirical scaling law identified by de Vries *et al* [250]. Agreement within the error bars is found for LFS sensors. On the other hand, the high field side (HFS) synthetic sensors measure a MP much weaker than expected from the de Vries scaling. This LFS–HFS asymmetry, which is clearly visible in figure 52, seems related to the fact that the various $m/1$ components interfere constructively on the LFS but destructively on the
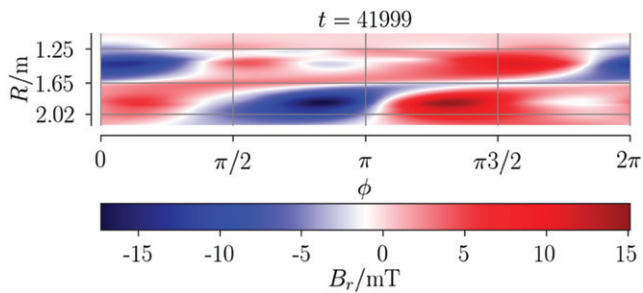
Figure 52. Perturbed radial magnetic field B_r at the midplane as a function of the toroidal angle ϕ and major radius R. The $q = 2$ surface (both on the HFS and LFS) and the magnetic axis are marked by grey lines. An HFS–LFS asymmetry clearly appears.
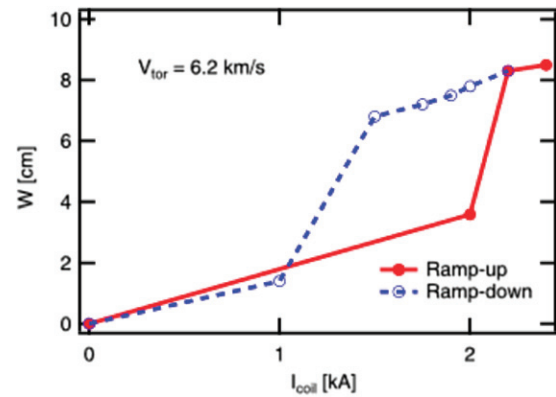

Figure 53. $2/1$ island width versus RMP coils current, showing the hysteresis effect: mode penetration occurs slightly above 2 kA in the RMP ramp-up phase, while mode expulsion occurs below 1.5 kA in the ramp-down phase. Reprinted from [255], with the permission of AIP Publishing.

HFS. A related observation is that the O-points of the various $m/1$ islands align on the LFS (a feature also observed experimentally [249]) while on the HFS, the alignment involves either the O-point or the X-point depending on the parity of m. These observations could help extract improved scaling laws in the future.

Furthermore, the locking of a slowly rotating magnetic island to the vacuum vessel of ASDEX Upgrade was demonstrated [248]. Scanning the vessel conductivity artificially shows that mode locking is fastest and most complete when $\tau_v \simeq m/\omega$, where τ_v is the vessel resistive time and ω is the mode frequency, as expected from theory [251].

6.1.2. Effect of impurities on island growth and relation to the Greenwald limit. Using the fluid impurity model (section 2.7), JOREK simulations show a strong growth of islands when the local radiation exceeds the Ohmic heating [252]. The transition to strong island growth occurs in a range of densities near the Greenwald limit. For a single impurity species, the critical density depends on the impurity fraction and the temperature, in contrast to the scaling of the Greenwald density limit. However, a mixture of impurity species is found to remove some of these dependencies, consistently with [253].

6.1.3. Tearing mode seeding via resonant magnetic perturbations (RMPs). In ASDEX Upgrade, experiments have been conducted using the RMP coils (the same as those used for ELM mitigation/suppression, see section 5.4) with a coil configuration optimized to produce $m/n = 2/1$ perturbations. Corresponding JOREK simulations show good qualitative agreement with the experiments and with analytical predictions regarding mode penetration [254]. In particular, thresholds for the mode penetration in the RMP amplitude and plasma rotation frequency were observed, as well as a fast transition between the shielded and penetrated states and a hysteresis of the island size and plasma rotation between the ramp-up and ramp-down of the RMP. The latter is illustrated in figure 53.

6.1.4. Tearing mode control with electron cyclotron current drive (ECCD). A useful tool to control TMs and help avoid disruptions is electron cyclotron current drive (ECCD). A fluid

closure reproducing the dominant Fisch–Boozer current generation mechanism in ECCD has been developed and validated against full Fokker–Planck calculations of ECCD [255]. This closure model is implemented in JOREK and the stabilizing influence of ECCD onto a TM has been demonstrated in simplified geometry [155].

6.2. Dynamics of disruptions triggered by massive material injection

This section summarizes work performed with JOREK on the TQ triggering mechanisms (section 6.2.1), the TQ dynamics and plasma current spike (section 6.2.2), the assimilation and mixing of injected material (section 6.2.3) and the radiation fraction and asymmetry (section 6.2.4). For work with other codes on these topics, see references [256–261] and references therein.

6.2.1. Thermal quench triggering mechanisms. Disruptions triggered by a MGI have been studied intensively with JOREK [74, 75, 78, 262]. An analysis of simulations of deuterium MGI in JET, illustrated by figures 54 and 55, suggests that the TQ is triggered through a current profile avalanche effect. The avalanche is started by the penetration of an MGI-driven cold front up to the $q = 2$ surface, at which point a large $m/n = 2/1$ TM is destabilized (first row of figure 54). The $2/1$ mode, whose growth is boosted by an MGI-induced helical cooling effect inside the $2/1$ island, flattens the toroidal current density (j_ϕ) profile around the $q = 2$ surface, which results in a steepening of the j_ϕ gradient more inward, as evident in the blue and red profiles in figure 55. When the steep j_ϕ gradient passes across the $q = 3/2$ surface, a $3/2$ TM is destabilized (second row of figure 54). The $3/2$ mode in turns flattens j_ϕ locally and propagates the steep j_ϕ gradient more inward, which destabilizes a $4/3$ TM (magenta curve in figure 55 and last row of figure 54). These combined modes generate global magnetic stochasticity and thereby provoke the TQ. It is interesting to note that, for these MGI simulations, the $n > 1$ modes seem to be destabilized via a current profile effect while in the study presented in section 6.1.1 they were destabilized by
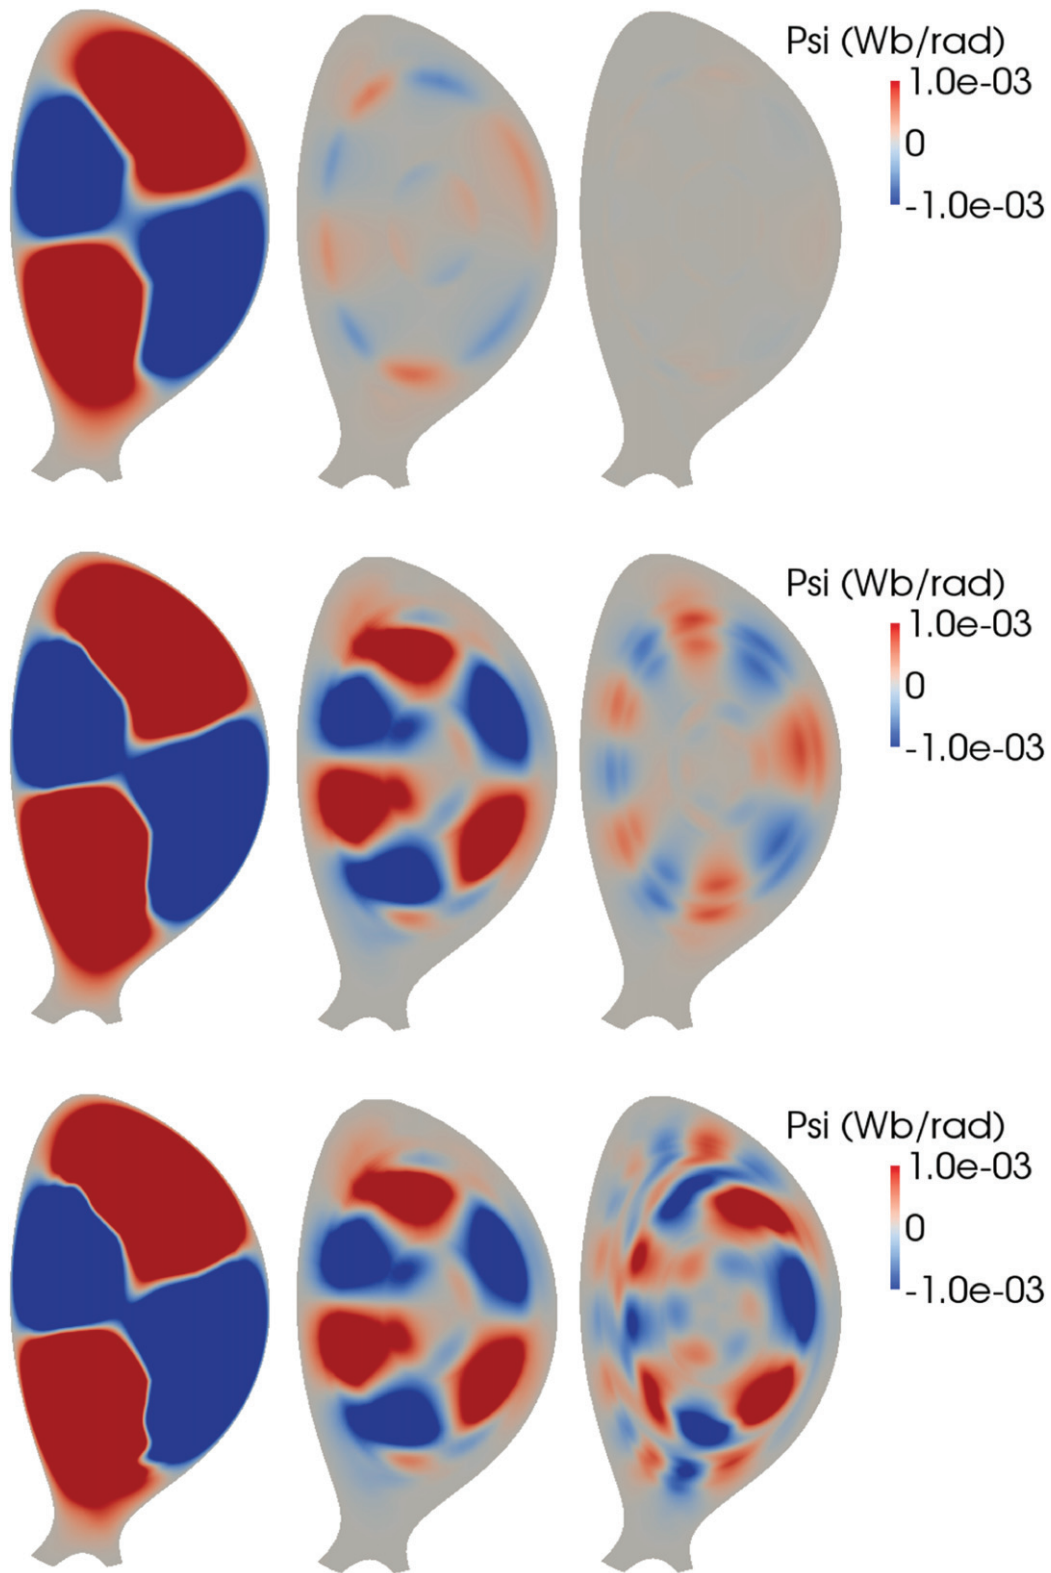
Figure 54. Poloidal cross-sections at the toroidal position of the MGI of (from left to right) the $n = 1$, $n = 2$ and $n = 3$ cosine component of the poloidal flux ψ at times (from top to bottom) $t = 4.1$ ms, $t = 5.1$ ms and $t = 5.7$ ms, for the same simulation of deuterium MGI in JET as in figure 55. The colour scale is the same for all plots (note the saturation for the $n = 1$ mode which has a large amplitude compared to the other modes). Reproduced from [79]. © IOP Publishing Ltd. All rights reserved.
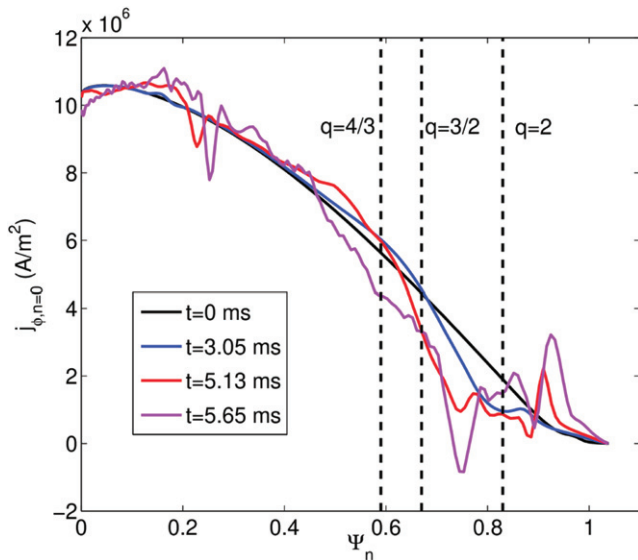
Figure 55. Toroidal current density profiles at the midplane (LFS) at different times for the same simulation of deuterium MGI in JET as in figure 54. The red and magenta profiles correspond to the last two rows of figure 54. Note that, for clarity, the position of rational surfaces is indicated (by vertical dashed lines) referring to their location at the beginning of the simulation. Reproduced from [79]. © IOP Publishing Ltd. All rights reserved.


Figure 56. Simulated and experimental time traces of the plasma current during a massive argon injection in JET (pulse 85 943). Vertical lines indicate the time of the Poincaré cross sections shown in figure 57.

mode beating. The reason for this difference remains to be clarified.

SPI is also being investigated in detail with JOREK. Simulations of deuterium SPI in JET [79, 87] and ASDEX Upgrade [263] have shown that pre-TQ dynamics can vary drastically depending on parameters. If shards travel across the plasma relatively fast compared to the current decay time in the SPI-cooled region, then the main MHD destabilization mechanisms are (1) the helical cooling on low order rational q surfaces (whereby shards may generate magnetic islands as they pass across these surfaces, possibly leading to a stochastization front progressing with the shards), and (2) for target plasmas with a central safety factor $q_0 < 1$, the excitation of the 1/1 internal kink mode when shards reach the $q = 1$ surface. This behaviour has been observed in JET deuterium SPI simulations [79, 87]. In the opposite limit of slow shards with respect to the current decay time in the SPI-cooled region, the dynamics resemble that described above for MGI: the current profile contracts due to the SPI-induced cooling of the edge, which destabilizes TMs in cascade from the edge to the core, leading to a stochastization front progressing faster than the shards. In this regime, the TQ is typically triggered when shards reach the $q = 2$ surface, but it can be triggered even before for large and slow pellets. This behaviour has been observed in ASDEX Upgrade deuterium SPI simulations [263] and is also typically observed for any tokamak when simulating the injection of shattered pellets containing impurities [80].

The key role of the ordering between the shards penetration time and the current decay time in the SPI-cooled region has a number of implications. First, it can be noticed that the
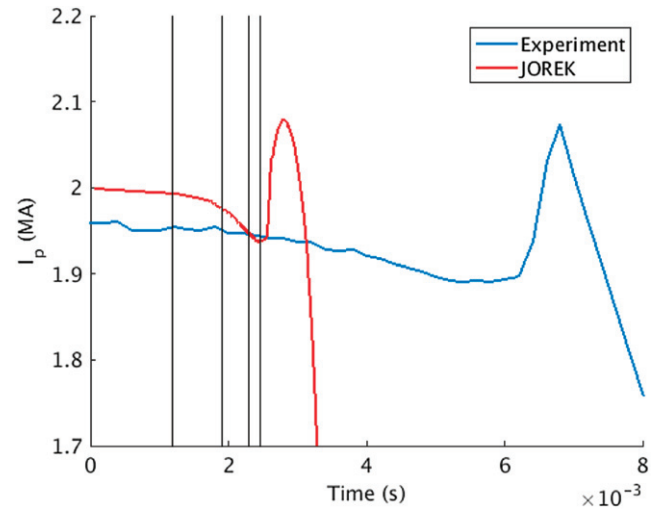
resistive current decay time, $\sim \mu_0 l^2 / \eta$, is proportional to the square of the length scale of interest l. Hence, considering that l is proportional to the machine size, the current decay time should grow like the machine size squared, while the shards penetration time, for a given velocity, increases linearly with machine size. This consideration is important for comparisons or extrapolations between machines of different sizes. It explains partly the different behaviour found in JET [79, 87] and ASDEX Upgrade [263] deuterium SPI simulations. Second, simulations in which the resistivity is artificially increased (which is often done for numerical reasons) should be considered with caution because the current decay time will be artificially shortened. To avoid this issue, present JOREK simulations use a realistic resistivity at low T_e with a cut-off above a certain T_e threshold. A third important remark is that the current decay time depends on the characteristics of the cooling, i.e. on (1) the timescale of the cooling and (2) the post-cooling temperature.

This last observation explains the large differences observed between simulations of pure deuterium versus impurity-containing MMI, and suggests that deuterium SPI may be used as a means to promptly and strongly dilute ITER plasmas without immediately triggering a TQ, thanks to a relatively high post-cooling temperature. Such a strategy may be instrumental in avoiding RE generation, motivating its investigation with JOREK [264]. It was found that in the absence of pre-existing islands and provided that the density of background impurities is low enough, the desired effect can indeed be obtained. Simulations are on their way to assess whether pre-existing magnetic islands can have a detrimental influence for realizing this strategy (first considering ASDEX Upgrade plasmas). It is assessed in particular whether the presence of the 2/1 island makes it harder to get SPI shards across the $q = 2$ rational surface without triggering a TQ [265].
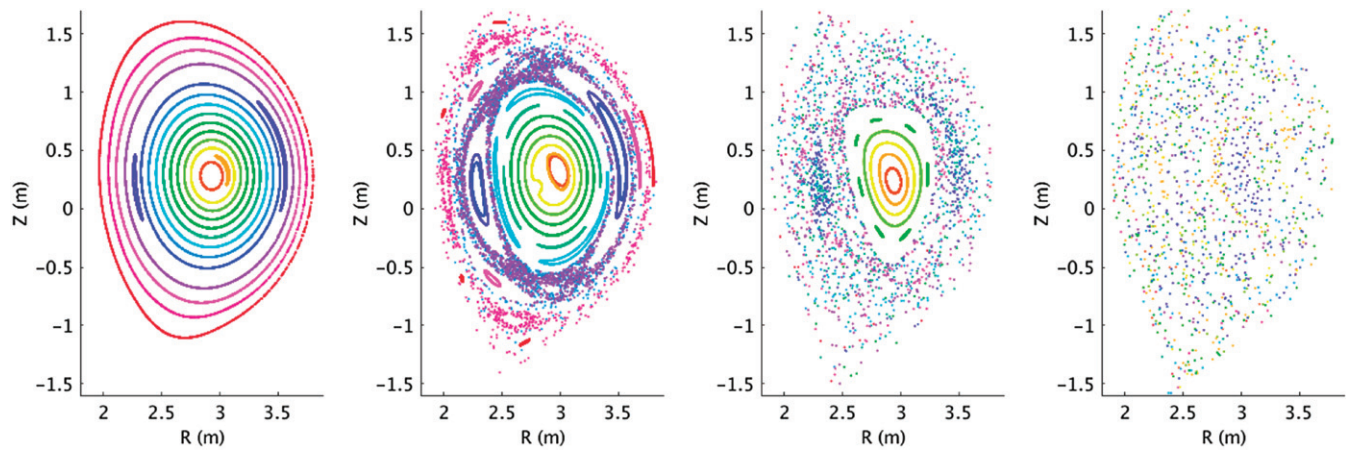
Figure 57. Poincaré cross sections at times 1.19 ms, 1.91 ms, 2.29 ms and 2.46 ms in a simulation of massive argon injection in JET (pulse 85 943). The corresponding times are indicated by vertical lines in figure 56.

6.2.2. Thermal quench dynamics and plasma current spike. The presence of an I_p spike is a robust experimental observation associated to the TQ, but a difficult one to reproduce quantitatively with 3D non-linear MHD simulations. I_p spikes in simulations have been reported in the past [78, 266] but with a significantly smaller amplitude than experimental spikes. However, in recent JOREK MGI simulations, pushing the parameters towards realistic experimental conditions and studying argon MGI in JET, an I_p spike of comparable magnitude to the experimental one has been obtained, as can be seen in figure 56. In these simulations, the mechanisms leading to the TQ are partly the same as described above for deuterium MGI but modes are more violently destabilized, leading to stronger stochasticity throughout the plasma (a series of Poincaré plots, at times indicated by vertical lines in figure 56, are shown in figure 57). The fact that modes are more destabilized is a result of a stronger cooling with argon than deuterium MGI and of the more realistic (i.e. lower) resistivity used in the simulation, which generates sharper skin currents and thus a more unstable current profile. Also, a seemingly critical feature associated to a large I_p spike [267] is a radiative cooling strong enough to persist near the O-point of the $2/1$ island, even as the island gets destroyed by magnetic stochasticity. This promotes a local collapse of the current density which drives the $2/1$ mode to a very large amplitude. According to theory [268], the I_p spike results from a relaxation of the current density profile at approximately fixed magnetic helicity. This relaxation is caused by the magnetic field stochastization and may be modelled, in a mean-field approach, by a hyper-resistivity term in the poloidal flux evolution equation [268]. Boozer predicts a relation between the mean-field hyper-resistivity and the field line stochastic diffusivity (equation (68) in [268]). JOREK mean-field simulations, i.e. axisymmetric simulations with an ad hoc hyper-resistivity term of the order of that predicted by Boozer, indeed match 3D simulations [269]. The I_p spike magnitude and the modes amplitude (which govern the field line stochastic diffusivity), thus appear to be strongly connected. In order to progress in the validation of JOREK simulations, ongoing work aims at comparing both the I_p spike magnitude (and more

generally the time evolution of I_p) and the modes amplitude with experimental data, in particular by using synthetic saddle coils in JET simulations [267].

Another important question related to RE generation, is that of the decay of MHD modes and magnetic stochasticity after the TQ. This point is actually also related to the I_p spike. Indeed, according to robust theoretical estimates done by Boozer, if the current profile relaxation during the TQ was complete, and assuming good magnetic helicity conservation, the I_p spike would be much larger than experimentally observed [268]. Thus, either helicity dissipation is non negligible, or the relaxation is incomplete. JOREK simulations seem to point to the latter explanation. Indeed, they usually display good helicity conservation during the TQ but a clear decay of stochasticity (starting from the core of the plasma) before the current profile has fully flattened, as visible for example in figures 11 and 12 of [263].

6.2.3. Assimilation and mixing of injected material. The question of the assimilation and mixing of the injected material is important, both for reducing localized heat fluxes and for avoiding REs. Indeed, the present ITER RE avoidance strategy relies on raising n_e by more than one order of magnitude throughout the plasma [270]. This density rise should be uniform, otherwise REs could be generated in low density regions.

It is generally observed in JOREK MGI and SPI simulations that the deposited material initially expands in the parallel direction at the speed of sound, as also seen in simulations of pellet ELM pacing (see section 5.2). After some time, the part of the cloud expanding in one direction may run into the part expanding in the other direction, leading to substantial viscous dissipation.

JOREK simulations of deuterium MGI and SPI in JET [79, 87] show that SPI is superior regarding material assimilation and mixing, as illustrated in figure 58, thanks to the deeper penetration of solid shards compared to gas. The same simulations also suggest, for target plasmas with q_0 substantially below 1 and thus with a large $q = 1$ radius, a key role of the $1/1$ internal kink mode for material mixing into the plasma core
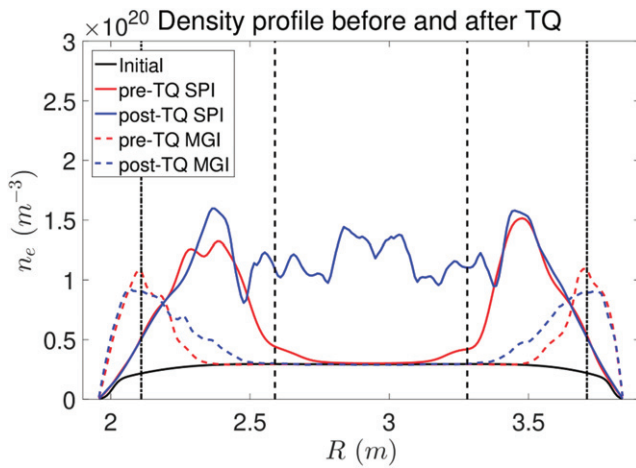
Figure 58. Electron density profile just before (red) and just after (blue) the TQ for simulations of deuterium SPI (plain) and deuterium MGI (dashed) in JET. The black plain and dashed vertical lines indicate the position of the $q = 2$ and $q = 1$ surfaces, respectively. Reproduced courtesy of IAEA. Figure from [88]. © 2018, ITER Organization.

during the TQ. In this respect, the position of the shards at the time of the internal kink mode crash appears critical: if shards are 'within reach' of the mode's flow structure, substantial material mixing into the core happens; otherwise, core mixing is poor. These observations are related to the 1/1 mode structure, which exhibits uniform and strong displacement within the $q = 1$ surface, but a weak displacement outside of that surface [271]. Simulations in which q_0 was elevated above 1 so as to remove the internal kink mode, also show poor core mixing early in the TQ. Recent impurity SPI simulations of ITER plasmas show similar results [80].

However, the absence of a 1/1 mode does not necessarily imply bad mixing. Indeed, simulations of deuterium SPI in ASDEX Upgrade [263], for which the target plasma has $q_0 > 1$, found good core mixing even if the TQ happens while shards are still far from the core (e.g. near the $q = 2$ surface). More precisely, as shown in figure 59, the density profile is usually hollow during and shortly after the TQ, but by the time flux surfaces reappear, the density profile has typically strongly flattened.

The more pronounced core mixing later into the TQ for the ASDEX Upgrade simulations is likely related to the much larger mode amplitudes and stronger magnetic stochasticity (as described in section 6.2.1), which lead to both stronger $\mathbf{E} \times \mathbf{B}$ flows and a larger radial transport by parallel flows. An important observation is that parallel flows are strongly driven during the TQ due to the heating of the relatively dense and cold region where the material has been deposited, by the heat flux coming from the core along stochastic field lines.

6.2.4. Radiated fraction and radiation asymmetry. A uniform deposition onto the plasma facing components (PFCs) of the pre-TQ thermal energy W_{th} is critical for efficient thermal load mitigation during ITER disruptions [247]. To achieve this, both a large radiated thermal energy fraction ($\geqslant 90\%$ for the baseline ITER scenario with $W_{th} = 350$ MJ) and a low radi-

ation asymmetry are required. It is usual to characterize the radiation asymmetry by a toroidal peaking factor (TPF) and a poloidal peaking factor (PPF). In this section, we discuss the TPF, which is defined as the maximal (over the toroidal angle) over mean radiated power per unit radian. The target for ITER is a TPF lower than 2. The PPF shall be the subject of future studies.

With SPI from a single toroidal location, the TPF may be above the ITER target during the pre-TQ and early TQ stages. An example of such behaviour is shown in figure 60. In this JET simulation, the injected fragments consist of pure neon and are flying along the poloidal plane. It is evident that the peak radiation power is located close to the toroidal location of the fragments, although there is some drift along the course of the injection. The TPF reaches its peak value, which is larger than 2, at the time of the TQ onset. It then gradually relaxes over the course of the TQ, until the radiation power becomes relatively toroidally uniform. The cause of such asymmetry can be traced back to the unrelaxed impurity density along the field lines, which is related to the finite ion sound velocity.

JOREK simulations show that multiple injections may significantly mitigate the TPF, as illustrated in figure 61. Here, we are comparing the toroidal radiation distribution of a single-SPI and a symmetric dual-SPI, both injecting 2.6×10^{22} neon atoms along with 2.1×10^{24} deuterium atoms into an ITER L-mode plasma. For the single-SPI, the injection is done at 0 radian, while for the dual-SPI injections are done at 0 and π radian. A significant TPF is again exhibited in the single-SPI case, even during the early stage of TQ at $t = 4.24$ ms. On the other hand, for the dual-SPI case, in spite of a strong asymmetry in the early pre-TQ stage, the asymmetry is strongly mitigated by the time of TQ at $t = 4.08$ ms.

6.3. Vertical displacement events and halo current dynamics

In the case where the vertical control of an elongated plasma is lost, the plasma column undergoes an axisymmetric instability referred to as a VDE. The loss of vertical control can be due to a large MHD perturbation or to a technical failure of the vertical control system. As the plasma drifts vertically towards the wall, large heat and electromagnetic loads are deposited into the PFCs and surrounding structures. The extrapolation of the magnitude and distribution of these loads to larger machines, in particular the 3D distribution, is still not well established. In this respect the simulation of these events with 3D MHD codes is crucial. For other non-linear MHD codes investigating VDEs, refer, e.g., to references [272–275] and references therein.

The simulation of VDEs requires free-boundary conditions for the magnetic field and a self-consistent evolution of currents in the structures surrounding the plasma. This is possible (including with 3D conducting structures) with the free-boundary extension of JOREK (see section 2.9). Based on this extension, a first benchmark of the growth rates for axisymmetric VDEs in a simplified ITER-like plasma had been carried out between JOREK and the CEDRES++ code in [168] revealing good quantitative agreement (within
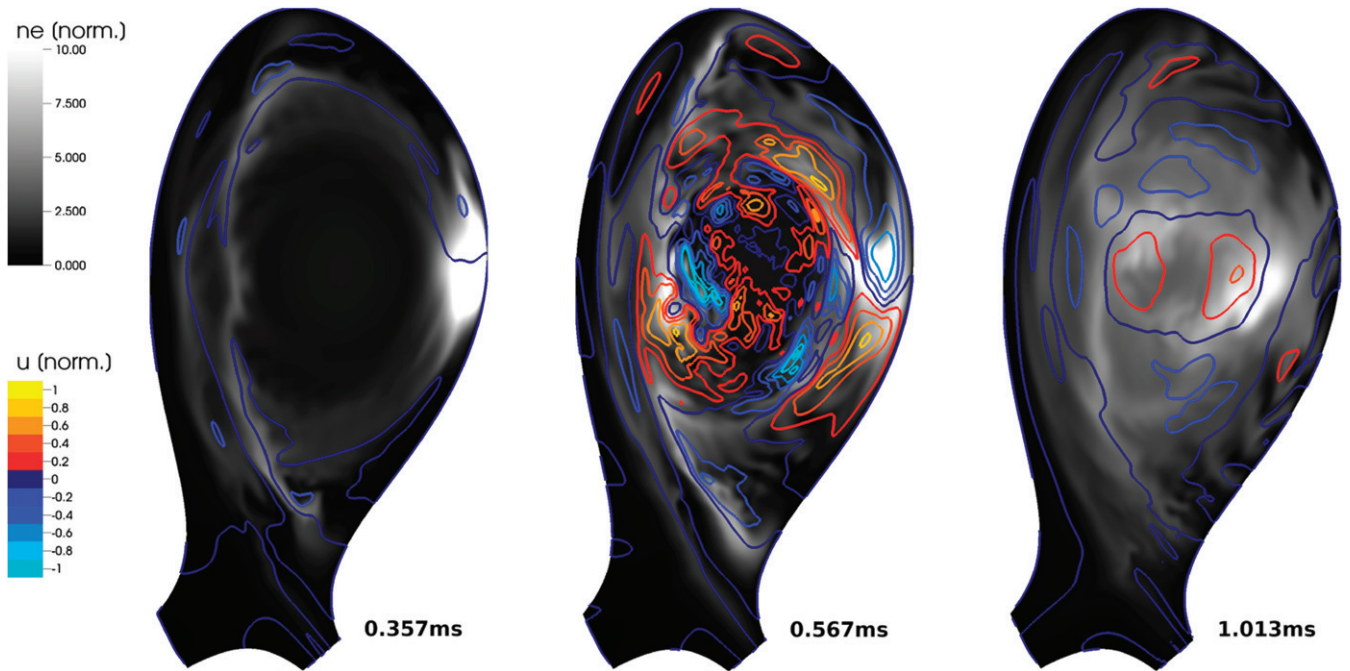
Figure 59. Colormap of the electron density distribution (black and white) with overlaid iso-contours of the stream function for the perpendicular plasma velocity (colours) during a simulation of deuterium SPI in ASDEX Upgrade. The 3 plots correspond, from left to right, to the pre-TQ, TQ, and early CQ phase. Reprinted from [264], with the permission of AIP Publishing.
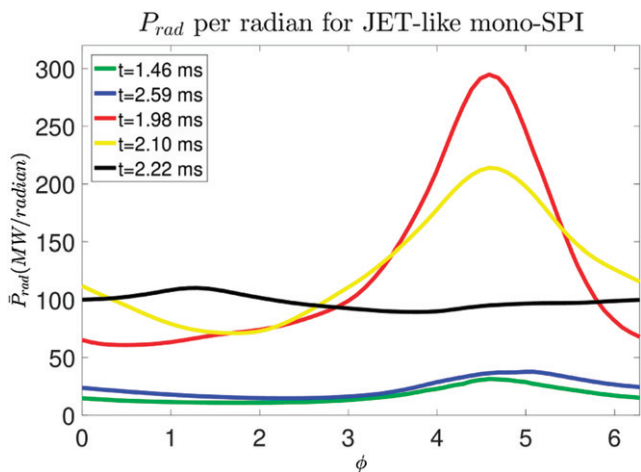
Figure 60. The relative and absolute radiation power within each poloidal plane for neon SPI into a JET L-mode plasma. The TQ occurs at $t = 1.98$ ms. The fragments are injected at a toroidal angle of 4.51 radian.

10%–20%) over a large range of wall resistivities (including fully realistic values). First demonstrations of 3D VDEs using JOREK were shown in [168, 276].

An important area of research connected to VDEs and disruptions is the halo current, i.e. the current that flows partly in the plasma SOL, closing its path via conducting structures in contact with the SOL [278–280]. Axisymmetric parametric scans for an ITER 15 MA upward VDE were performed in [277] by fixing the halo region temperature and width. Scans in the plasma resistivity showed that when the CQ time is much smaller than the decay time of the wall currents ($\tau_{CQ} \ll \tau_w$),

the vertical position is a monotonic function of the plasma current $Z(I_p)$ as predicted by Kiramov's wire model [281]. In this limit the shape of the $Z(I_p)$ curve does not depend on specific time scales but only on the shape of the toroidal current profile. The implication is that ITER plasmas will transition from an X-point to a limiter configuration at a large total current (~ 10 MA) regardless of the mitigation scheme used, which could have important consequences regarding the wall damage caused by a vertically unstable RE beam, if RE mitigation by SPI does not achieve its goal in ITER.

The latter scans show that the expected halo current fraction (HF) depends strongly on the τ_{CQ}/τ_w as shown in figure 62. A minimum HF of $\sim 10\%$ is found at $\tau_{CQ}/\tau_w \ll 1$ and maximum fractions of $\sim 50\%$ have been found in the limit $\tau_{CQ}/\tau_w \gg 1$, similar to the maximum values obtained with the DINA code [247]. In the fast CQ limit, the maximum HF is below 10% regardless of halo width and temperature assumptions. Mitigated disruptions in ITER will target CQ times in the range of 50–150 ms (to be compared to $\tau_w = 0.5$ s). The corresponding predicted HF range is 10%–20%. This suggests that the ITER DMS could be used to reduce the halo current electromagnetic loads by more than a factor 2 with respect to their maximum possible value.

3D simulations for an ITER 15 MA plasma were conducted in the fast CQ limit [92]. The TPF of the poloidal halo currents is shown in figure 63 for 3 different values of the plasma resistivity leading to different CQ times. It can be seen that the TPF (given by an $n = 1$ kink mode) is very low in the fast CQ limit (while in experiments, the TPF can reach values of 4). Moreover the maximum TPF does not seem to depend strongly on the CQ time if $\tau_{CQ}/\tau_w \ll 1$. Note however that CQ times in these simulations are below 50 ms, which is not allowed in
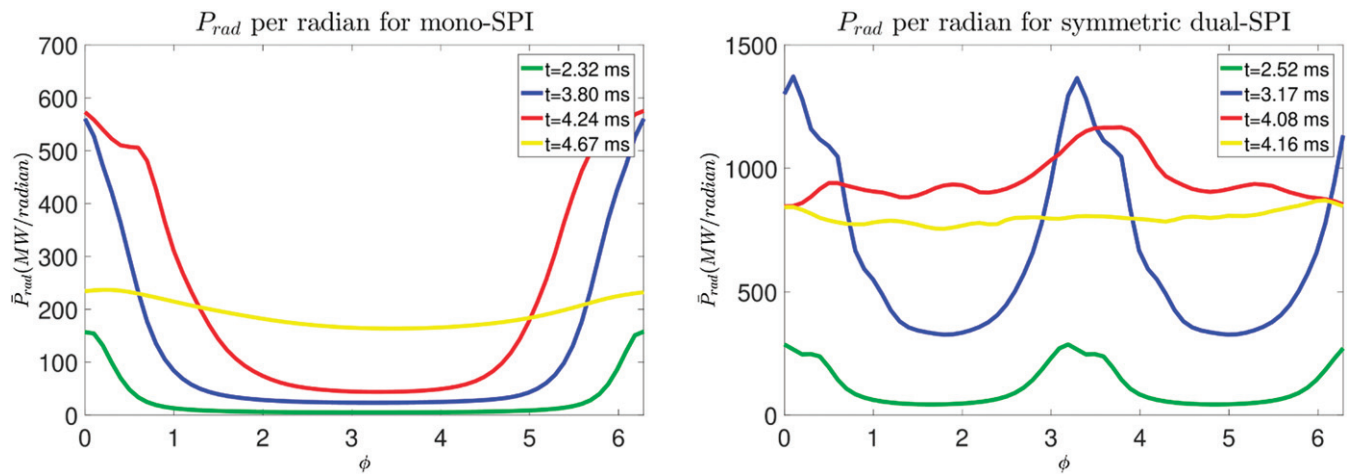
Figure 61. Comparison of radiation asymmetries between mixed neon-deuterium single-SPI (left) and dual-SPI (right) into an ITER L-mode plasma. Reproduced courtesy of IAEA. Figure from [81]. Copyright 2020 IAEA.
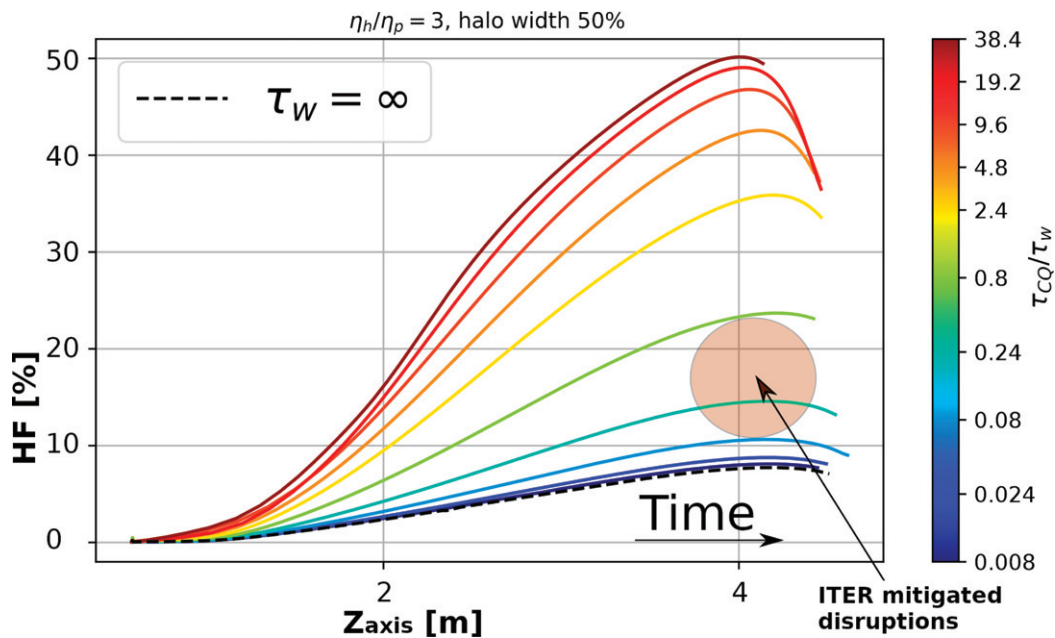

Figure 62. Halo fraction versus vertical position of the magnetic axis for different τ_{CQ}/τ_{w} ratios [277]. The halo current fraction is defined as $HF \equiv (1/I_{p,0}) \int |\mathbf{J} \cdot \mathbf{n}|/2 \, dS_{wall}$.

ITER due to eddy current constraints on the blanket modules [247]. These results are therefore presented only for their physical interest and future work will be devoted to realistic ITER predictions.

The change of the edge current density and edge safety factor during a VDE when the plasma moves into the PFCs is discussed in [282]. Analytical theory and JOREK non-linear simulations show excellent agreement. When currents are lost in the scraped-off region of the plasma, a significant fraction of these currents is re-induced in the edge hot core region of the plasma. The same mechanisms also apply in case of a loss of the plasma edge current due to a cooling of the plasma edge caused by an MMI.

Predictive simulations of VDEs must include several effects that are relevant for the evolution of the density and of the temperature in the SOL region. The energy balance given by perpendicular and parallel transport, Ohmic heating, radiation, ionisation and other effects determines the plasma temperature, which in turn has a large impact on the current density and the VDE dynamics through the plasma resistivity. In [67], a complete set of sheath boundary conditions was taken into account together with a neutral fluid model and realistic physical parameters, e.g., Spitzer resistivity and Spitzer–Härm parallel conduction. It was found that it is particularly important to couple the electrical current density to the particle density and temperature through a boundary condition that limits the maximum current density to the ion saturation current. Otherwise large currents can be induced in regions of very low particle density, leading to non-physical dynamics. Although the effect of impurity radiation was not included, these
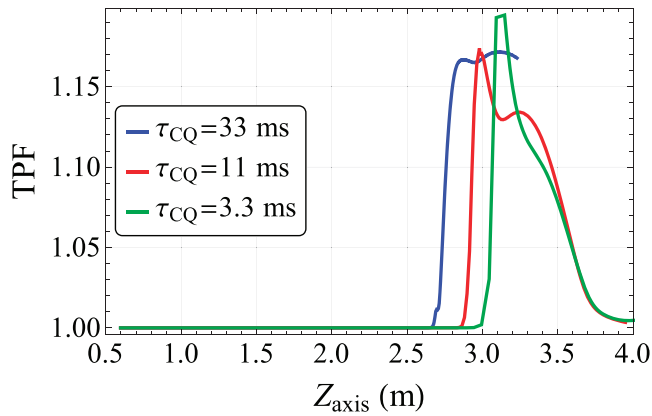
Figure 63. TPF of the halo current as a function of the vertical position of the magnetic axis for an ITER 15 MA upward VDE [92].

simulations are the first to provide a fully self-consistent evolution of halo currents.

6.4. Runaway electron physics

6.4.1. Test electron dynamics during the thermal quench. RE generation during disruptions is a major concern for ITER [283]. A key question in this area is that of fast electron losses along stochastic field lines during a disruption. To study this question, a relativistic test particle tracer (with 2 options: full orbit or guiding centre) has been implemented in JOREK (see section 2.10). It has first been applied to study electron losses at approximately fixed energy in the electromagnetic fields of MGI-triggered disruptions (see section 6.2) [101, 102]. Substantial electron losses were found during the strongly stochastic TQ phase, which increased with the electron energy until reaching a saturation as the electron velocity approaches the speed of light. However, a few tens of % of electrons at typical pre-TQ thermal energies ($\sim$1 keV) were found to remain in the plasma by the time flux surfaces start reforming after the TQ. This is by far enough to give birth to a large RE beam, but this number should be taken with caution since the magnetic field stochasticity may have been underestimated in these simulations, as suggested by the under-predicted plasma current spike. In a second study, the full dynamics of test electrons, including acceleration or braking by the parallel electric field and by collisions, were investigated in the same JOREK simulations [102, 284]. This revealed that the very large ($\sim$1 kV m^{-1}) parallel electric field fluctuations during the TQ can have a key role in RE generation. These fluctuations indeed rapidly broaden the electron energy distribution, with a typical energy reaching a few tens of keV after a few tens of μs, as can be seen in figure 64. This makes electrons less collisional and therefore promotes RE generation. This is a new effect, not captured by typical lower-dimensional studies where only the axisymmetric component of the electric field is considered. Finally, the magnetic field from a JOREK disruption simulation has been used in a recent test electron study with the ASCOT code [175].

6.4.2. Runaway electron fluid model applied to VDEs and RE beam termination. Complementary to the passive test particles approach, an RE fluid model has been implemented, which allows to describe the full non-linear interaction of the REs with the MHD activity [110]. While various applications of the model are on their way, some studies have already been performed. Similar approaches are followed by other codes, see e.g., references [167, 285–287].

Axisymmetric and non-axisymmetric VDEs have been studied with and without the presence of REs [288], showing that the vertical motion of the plasma can be slowed down by the presence of REs if the decay of the plasma current is influencing the dynamics of the vertical instability significantly. It was also demonstrated that the RE current, peaking off-axis in this case, may lead to the destabilization of a 1/1 internal kink instability in a plasma that would otherwise have a dominant 2/1 helical perturbation, as shown in figure 65. Such effects may have important consequences for the termination of RE beams.

In a recent study, the benign termination of an RE in a JET experiment has been simulated [289, 290]. Based on experimental data, a plateau-phase RE beam with a hollow current density profile in a low density background plasma with negligible content of high-Z impurities is considered. It is found that an $n = 1$ double TM associated with the presence of two $q = 4$ surfaces grow rapidly, with subsequent non-linear mode interaction causing global magnetic stochastization. This leads to the loss of REs on a timescale of $\leqslant 100\,\mu$s. Poincaré plots as well as the RE number density are shown in figure 66. The observed behaviour is in very good agreement with the experiment as it shows comparable mode structures and time scales. Simulations also indicate a significant toroidal variation in the RE flux on the wall dominated by an $n = 1$ structure and a poloidally broad RE deposition zone, partly explaining why no material damage is observed experimentally.

6.5. Outlook

Concerning current and future work related to disruptions, like mentioned in section 2.15, the development of a more realistic impurity model going beyond the coronal equilibrium assumption, is ongoing to improve the accuracy of MMI simulations.

A validation effort on MMI modelling is underway, involving in particular the simulation of MGI and SPI experiments in JET. Synthetic diagnostics such as interferometry, bolometry and saddle loops have been implemented for this purpose.

As described above, a growing effort is also devoted to testing disruption mitigation strategies for ITER. In particular, the study on the possibility of promptly diluting ITER plasmas with pure deuterium SPI, described in section 6.2.1, is being refined, investigating e.g. the effect of pre-existing islands. Also, the work aimed at optimizing thermal load mitigation by radiation, described in section 6.2.4, is being pursued, considering more ITER scenarios and refining the treatment of
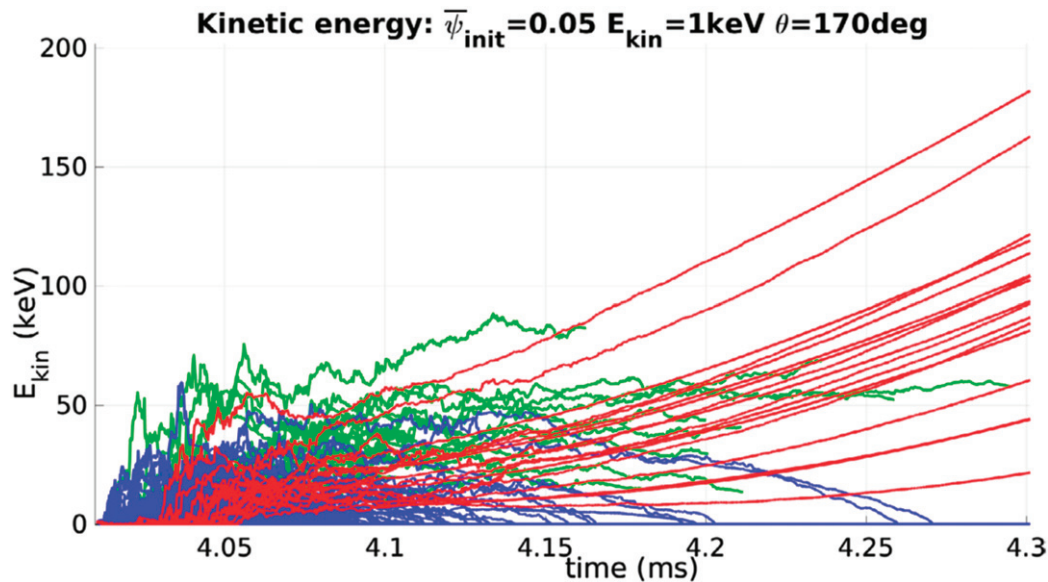
Figure 64. Early evolution of the kinetic energy of a set of electrons initialized in the plasma core with a kinetic energy of 1 keV. Green lines correspond to lost electrons while red (resp. blue) lines correspond to electrons which remain confined and have a final energy above (resp. below) 1 MeV. Reproduced courtesy of IAEA. Figure from [285]. © 2018 EURATOM.
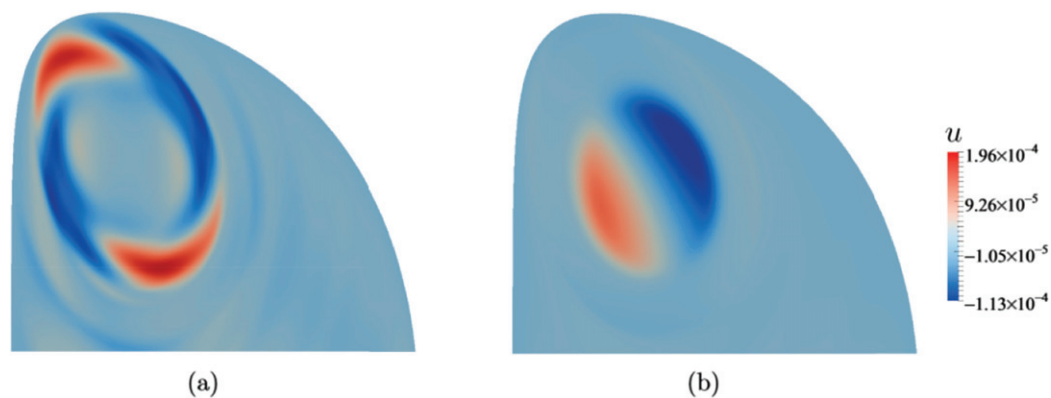
Figure 65. Structure in the perturbed (normalized) electric potential u during the non-linear phase of the mode growth. (*a*) Without REs, the $n = 2$ component dominates while (*b*) with REs an $n = 1$ mode is developing due to the different q-profile (1/1 internal kink). Reprinted figure with permission from [111], Copyright 2019 by the American Physical Society.

impurities. The RE fluid model will be applied to further test RE avoidance or mitigation schemes for ITER.

In the field of VDEs, current and future work involves experimental validation, implementation of sheath boundary conditions, a free-boundary extension for full MHD and a coupling with 3D volumetric wall codes.

7. Further applications

7.1. ITGs

As described in section 2.13, an electrostatic model for ITG turbulence has been recently implemented. The fluid model and first benchmark results were shown in reference [116]. The effect of magnetic shear onto the characteristics of global ITG modes are investigated in reference [117] in global simulations in simplified circular geometry similar to [118]. Typical up-down asymmetrical eigenmode structure with non zero

ballooning angle is typically observed in the regimes of relatively large magnetic shear due to the increased coupling between rational surfaces [117]). In the low shear regime, the unstable eigenmodes become narrowly localized on the corresponding rational surfaces and exhibit no up-down asymmetry. The structure of the generated mean poloidal flow via Reynolds stress is investigated in more detail in reference [117].

The first benchmark results for linear ITG modes of the full-f electrostatic kinetic JOREK model with full orbit ions and adiabatic electrons are shown in figure 67. For these cases, the full orbit of 10^9 ions are traced with a time step of 5×10^{-9} s in a polar flux surface-aligned finite element grid of 101 radial and 256 poloidal elements. The full orbit ITG growth rates from JOREK are in good agreement with the gyrokinetic results from GENE, XGC and ORB5 presented in reference [119]. Since the kinetic model uses the same grids and
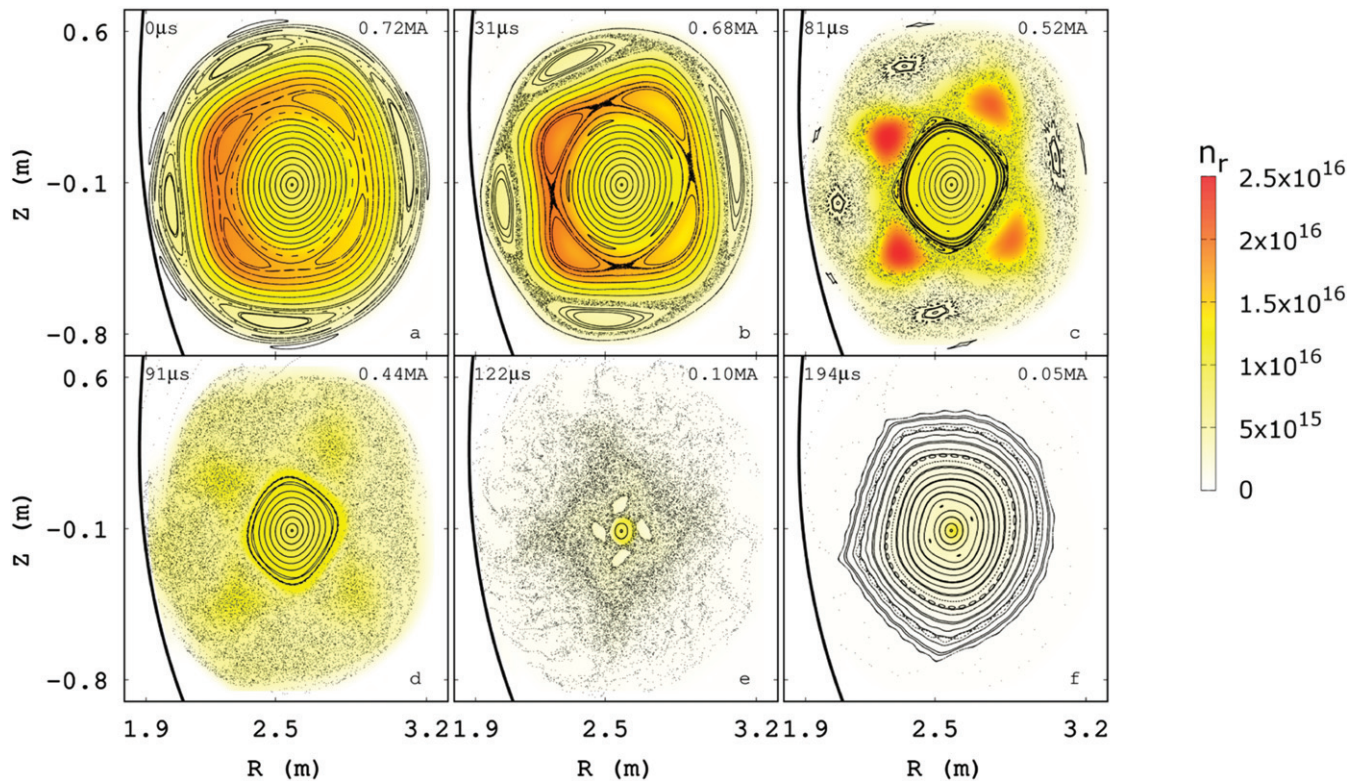
Figure 66. Poincaré plots along with the RE number density (background colour) during the MHD activity leading to the stochastization of the magnetic field and subsequent termination of the RE beam. Remaining RE current is also shown in each panel. Reproduced from [292]. © IOP Publishing Ltd. All rights reserved.
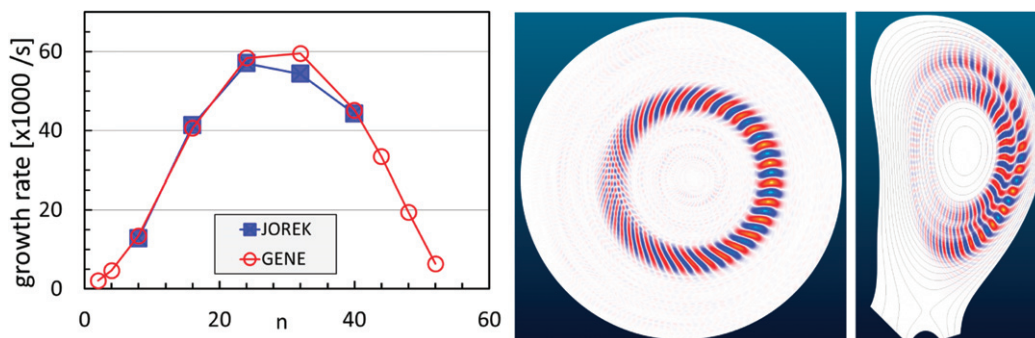
Figure 67. (Left) Growth rate of ITG modes as a function of the toroidal mode number from the JOREK kinetic model compared to the GENE results from the benchmark [119]. (Middle) The ITG mode structure of the potential for the $n = 24$ case. (Right) Example of the mode structure of an $n = 20$ ITG mode in COMPASS X-point geometry L-mode plasma.

equilibria as the fluid models in JOREK, the kinetic model can also be applied in X-point geometry, including open field lines. As an illustration of this, figure 67 (right) shows an $n = 20$ ITG mode in COMPASS X-point geometry, in the non-linear phase.

7.2. Toroidal Alfvén eigenmodes and fast particles

The excitation of TAE modes with an external antenna was studied in JET X-point plasmas was studied in references [291, 292] using the freeboundary extension described in section 2.9. The paper addresses the question why, in JET experiments, TAE modes are more difficult to excite in

X-point geometry as opposed to limiter geometry. The JOREK simulations show the same behaviour as observed experimentally and point to the importance of the SOL with open field lines with a wider SOL leading to a less efficient antenna excitation.

The particles framework (section 2.10) includes the interaction of fast particles on the MHD fluid. The fast particle driven TAE instability has benchmarked against the well-established ITPA TAE benchmark as described in [293]. In the JOREK simulations, the pressure coupling scheme was used and the number of particles was 10^8. Figure 68 shows a good agreement in the comparison of the JOREK results with several
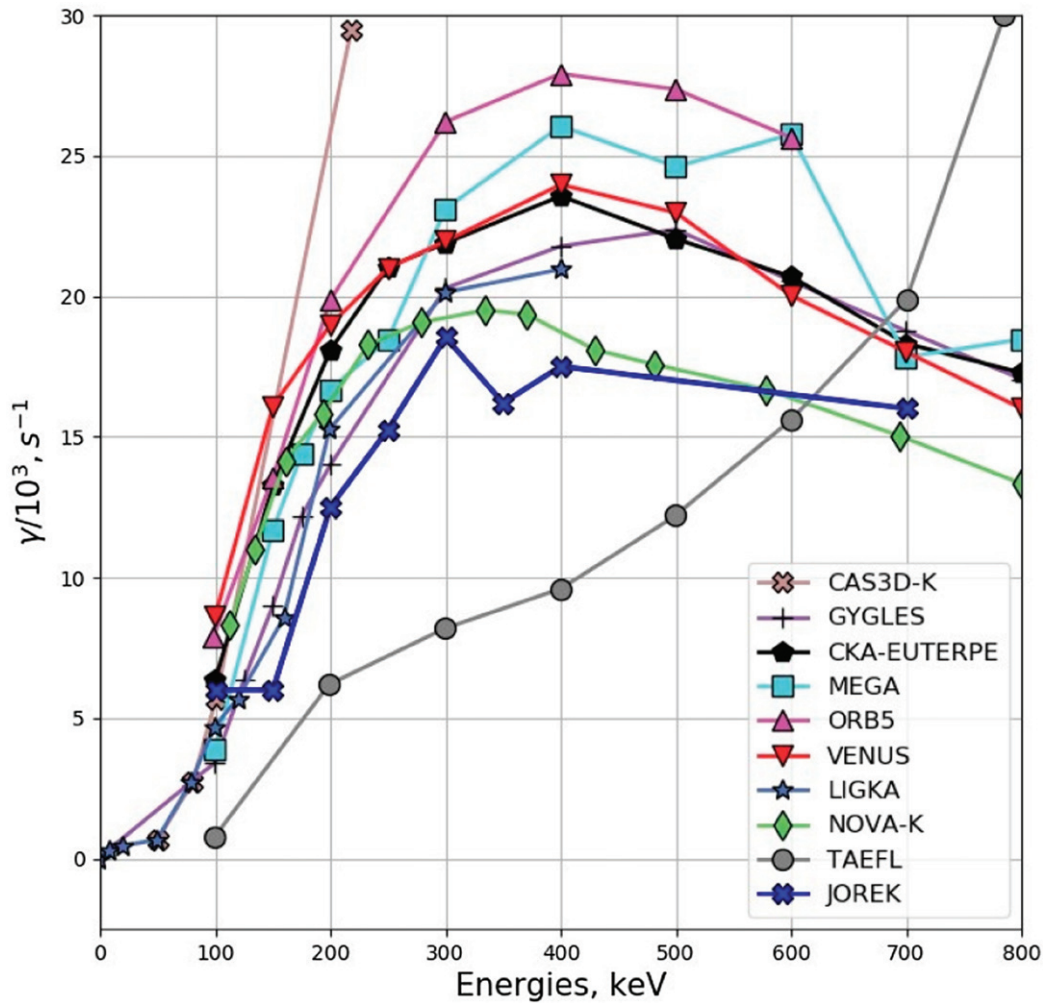
Figure 68. The linear growth rate of the $n = 6$ TAE mode as a function of the temperature of the fast particle population, comparing JOREK results to other codes (including FLR effects). The data is taken from reference [293] with the JOREK results from reference [294] added.

other codes on the $n = 6$ TAE mode growth rate as a function of the fast particle temperature [294]. Note that JOREK is the only code in this benchmark using full orbit fast ions.

7.3. 3D configurations

3D plasma configurations are becoming an important research area for JOREK, in particular quasi-axisymmetric (QA) stellarators are of interest. A first approach is to simulate an axisymmetric configuration with properties as close as possible to the true QA configuration. Linearly, such an approach has been followed in reference [295] and is presently improved by including 'virtual coil currents' in the simulations, which provide an externally driven rotational transform like it is present in the QA-stellarator also for the axisymmetric simulation [296–298]. This allows to incorporate both the influence of the external rotational transform onto non-axisymmetric modes and the stabilizing effect onto VDEs. In parallel to this axisymmetric approach to the QA stellarator, an extension of JOREK to 3D configurations is presently ongoing. As a first step, a hierarchy of reduced and full MHD models has been derived in a form suitable for stellarator devices [70, 71]; implementation is presently on the way. The extension of the

code to 3D grids taking place in parallel is briefly mentioned sections 2.15 and 3.8. First simulations in simple stellarator geometry are expected to become possible very soon.

8. Conclusions and summary

A comprehensive summary of the JOREK simulation code was given covering the available physics models, the numerical methods employed and the verification, as well as the broad range of applications to magnetically confined fusion plasmas. An outlook was provided in each section onto further developments of the code and future applications.

It was shown that JOREK provides a framework which contains a number of different physics models ranging from reduced and full MHD models to a fully kinetic treatment of the plasma (so far electrostatic). Various extensions are available including separate electron and ion temperatures, diamagnetic drift, neoclassical effects, fluid and kinetic neutrals, a RE fluid and test particle model, fluid and kinetic impurities, pellets, free boundary and conducting structures. While the single fluid reduced and full MHD models are energy conserving on the equation level, errors can arise from gyro-viscous cancelation, temporal discretization and too low toroidal resolution.

Diagnostics running automatically for each simulation allow to confirm that energy is conserved reasonably well in practice. Momentum conservation is exact on the equation level for the full MHD model, but not for the reduced MHD model. The error has a higher order as seen from analysis in reference [70] and confirmed by the good linear and non-linear agreement between reduced and full MHD models shown in direct comparisons.

The spatial discretization is based on a 2D G^1 continuous finite element formulation combined with a toroidal Fourier decomposition allowing to accurately align to flux surfaces improving numerical accuracy and to extend the computational domain across the separatrix up to divertor and plasma facing components. An extension of the axisymmetric grid to 3D stellarator configurations is under development. The robust fully implicit time advance allows to use large time steps where the physics processes allow for it. An iterative solver with a physics based preconditioner is applied to the large sparse matrix system, for which various options and library interfaces exist and further improvements are on their way. The code is very actively developed in an international community with automatic regression tests and code reviewing.

A large number of verification activities have been carried out over the years, while only selected ones could be shown in this article. This covers basic convergence properties of the discretization, tests for highly anisotropic heat transport, a verification of energy conservation, linear and non-linear benchmarks on core and edge instabilities, comparisons on the SOL models, benchmarks for the RE fluid model, a variety of simpler and very advanced benchmarks for the free boundary extension, and detailed tests of the kinetic particle framework.

Regarding pedestal/edge simulations, a variety of results regarding natural ELMs was shown demonstrating that the JOREK code is able to reproduce key experimental observations qualitatively and even quantitatively. This is in particular the case for the divertor heat fluency scaling that is recovered reasonably well and for the explosive onset of type-I ELM crashes that has recently been reproduced for the first time in simulations of full ELM cycles. Regarding the control of ELMs, pellet ELM triggering has been studied extensively and, for instance, the lag-time experimentally observed in the ELM cycle has been reproduced recently in simulations with fully realistic parameters and flows. ELM triggering via vertical magnetic kicks has been demonstrated as well. The penetration of error fields into the plasma as well as the mitigation and suppression of ELMs via RMP fields was investigated in detail including realistic $E \times B$ and diamagnetic flows. ELM free QH-mode regimes were obtained and the EHO was explained by a saturated KPM. Using a neutral fluid model, the ELM burn-through in detached conditions was demonstrated in non-linear simulations and a more precise description of the SOL/divertor processes is on the way including a kinetic treatment of impurities and neutrals.

Regarding disruptions, results were presented on predisruption physics, the dynamics of MMI triggered disruptions, VDEs, halo currents, and REs. Most prominently,

JOREK is able to capture, at least in a qualitative sense, experimentally observed features of disruptions like the triggering of a TQ via the destabilization of a 2/1 TM and the ensuing non-linear dynamics, and the I_p spike. It can also describe accurately 3D VDEs and halo current dynamics in spite of using a reduced MHD model, as benchmarks with full MHD codes have shown. While work is ongoing towards quantitative validation of disruption simulations, JOREK is already being used to test and optimize disruption mitigation strategies for ITER, and these efforts will intensify in the coming years.

Further applications of JOREK include ITG turbulence simulations, the interaction of TAE modes with fast particles, and the ongoing extension towards stellarators.

Acknowledgments

Appendix A. Coordinate systems

Cylindrical coordinates. As already mentioned in section 2.2, the basic cylindrical coordinate system $(u^1, u^2, u^3) = (R, Z, \phi)$ of JOREK is given by $x = R \cos \phi$, $y = -R \sin \phi$, and $z = Z$, where (x, y, z) denotes Cartesian coordinates (figure 1). The tokamak coordinate convention number COCOS [43] used by JOREK is 8. This implies the COCOS coefficients $e_{Bp} = 0$, $\sigma_{Bp} = -1$, $\sigma_{R\phi Z} = -1$, $\sigma_{\rho\theta\phi} = 1$. The covariant basis vectors $\mathbf{a}_\alpha = \partial \mathbf{X}/\partial u^\alpha$ are given by

$$\mathbf{a}_1 = \begin{pmatrix} \cos \phi \\ -\sin \phi \\ 0 \end{pmatrix}, \quad \mathbf{a}_2 = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}, \quad \mathbf{a}_3 = \begin{pmatrix} -R \sin \phi \\ -R \cos \phi \\ 0 \end{pmatrix}.$$

We use both 1, 2, 3 and R, Z, ϕ synonymously as sub- or superscripts in our notation to identify the co- and contravariant components such that, e.g., $\mathbf{a}_\phi \equiv \mathbf{e}_3$. The cross products between these basis vectors are $\mathbf{a}_1 \times \mathbf{a}_2 = \mathbf{a}_3/R$, $\mathbf{a}_1 \times \mathbf{a}_3 = -R\mathbf{a}_2$, $\mathbf{a}_2 \times \mathbf{a}_3 = R\mathbf{a}_1$, and $\mathbf{a}_\alpha \times \mathbf{a}_\alpha = 0$ as well as $\mathbf{a}_\alpha \times \mathbf{a}_\beta = -\mathbf{a}_\beta \times \mathbf{a}_\alpha$. The contravariant basis vectors are given by $\mathbf{a}^1 = \mathbf{a}_1$, $\mathbf{a}^2 = \mathbf{a}_2$, $\mathbf{a}^3 = \mathbf{a}_3/R^2$. Of course, $\mathbf{a}_\alpha \cdot \mathbf{a}^\beta = \delta_\alpha^\beta$ and $\mathbf{a}^1 = \nabla R$, $\mathbf{a}^2 = \nabla Z$, $\mathbf{a}^3 = \nabla \phi$. Furthermore, $\mathbf{a}_1 = J\nabla Z \times \nabla \phi$, $\mathbf{a}_2 = J\nabla \phi \times \nabla R$, $\mathbf{a}_3 = J\nabla R \times \nabla Z$. And $\mathbf{a}^1 \times \mathbf{a}^2 = R\mathbf{a}^3$, $\mathbf{a}^1 \times \mathbf{a}^3 = -\mathbf{a}^2/R$, $\mathbf{a}^2 \times \mathbf{a}^3 = \mathbf{a}^1/R$. Here the Jacobian is $J = \mathbf{a}_1 \cdot (\mathbf{a}_2 \times \mathbf{a}_3) = R$. Normalized basis vectors are given by $\mathbf{e}_1 = \mathbf{e}^1 = \mathbf{a}_1 = \mathbf{a}^1$, $\mathbf{e}_2 = \mathbf{e}^2 = \mathbf{a}_2 = \mathbf{a}^2$, $\mathbf{e}_3 = \mathbf{e}^3 = \mathbf{a}_3/R = R\mathbf{a}^3$. The co- and contravariant metric tensors are given by $g_{\alpha\beta} = \mathbf{a}_\alpha \cdot \mathbf{a}_\beta = \mathrm{diag}(1, 1, R^2)$, $g^{\alpha\beta} = \mathbf{a}^\alpha \cdot \mathbf{a}^\beta = \mathrm{diag}(1, 1, 1/R^2)$. The determinant of the covariant metric tensor is $g \equiv J^2 = \det(g_{\alpha\beta}) = R^2$. Differential operators are given in the cylindrical coordinates of JOREK by

$$\nabla U = \partial_1 U \mathbf{a}^1 + \partial_2 U \mathbf{a}^2 + \partial_3 U \mathbf{a}^3$$

$$\nabla_{\mathrm{pol}} U = \partial_1 U \mathbf{a}^1 + \partial_2 U \mathbf{a}^2$$

$$\nabla \cdot \mathbf{V} = \frac{1}{R}\partial_1(RV^1) + \partial_2 V^2 + \partial_3 V^3$$

$$\nabla \times \mathbf{V} = \frac{1}{R}(\partial_2 V_3 - \partial_3 V_2)\mathbf{a}_1 + \frac{1}{R}(\partial_3 V_1 - \partial_1 V_3)\mathbf{a}_2$$
$$+ \frac{1}{R}(\partial_1 V_2 - \partial_2 V_1)\mathbf{a}_3$$

$$\Delta U = \nabla \cdot \nabla U = \frac{1}{R}\partial_1(R\partial_1 U) + \partial_{2,2} U + \frac{1}{R^2}\partial_{3,3} U$$

$$\Delta_{\mathrm{pol}} U = \nabla \cdot \nabla_{\mathrm{pol}} U = \frac{1}{R}\partial_1(R\partial_1 U) + \partial_{2,2} U$$

$$\Delta^* U = R^2 \nabla \cdot \left(\frac{1}{R^2}\nabla_{\mathrm{pol}} U\right) = R\partial_1\left(\frac{1}{R}\partial_1 U\right) + \partial_{2,2} U$$

$$[A, B] = \mathbf{e}_3 \cdot (\nabla A \times \nabla B) = \partial_1 A \partial_2 B - \partial_2 A \partial_1 B$$

The Christoffel symbols defined by

$$\mathbf{A} \cdot \nabla \mathbf{B} = A^i \partial_i(B^j a_j) = A^i(\partial B^j)a_j + A^i B^j(\partial_i a_j)$$
$$= A^i(\partial_i B^j + \Gamma_{ik}^j B^k)a_j$$

$$(\mathbf{A} \cdot \nabla \mathbf{B})^j = A^i(\partial_i B^j + \Gamma_{ik}^j B^k)$$

are given by the following expressions in the JOREK cylindrical coordinate system:

$$\partial_3 \mathbf{a}^1 = R\nabla \phi = R\mathbf{a}^3, \qquad \partial_3 \mathbf{a}^3 = -\frac{1}{R}\mathbf{a}^1, \qquad \partial_1 \mathbf{a}^3 = -\frac{1}{R}\mathbf{a}^3.$$

Element local coordinates. Inside each of the quadrangular finite elements (see section 3.1), a local coordinate system (s, t, ϕ) is defined with s and t taking values in the range $[0, 1]$. Both s and t are orthogonal with respect to ϕ, but not with respect to each other. The local coordinates are defined by the cylindrical coordinates $R(s, t)$ and $Z(s, t)$ being expressed in terms of the local coordinates in the same finite element basis as the physical variables. When the Poisson bracket $[a, b]$ is expressed in terms of the element–local coordinates, we get $[a, b] = (a_{,s}b_{,t} - a_{,t}b_{,s})/J_2$ where $J_2 = R_{,s}Z_{,t} - R_{,t}Z_{,s}$. Note, that the singularity at the grid centre, where $J_2 = 0$ does not break the code, since the integration is carried out on the Gaussian integration points, where $J_2 \neq 0$.

While this mapping does not have a ϕ dependency at present, an extension of the code to 3D grids is presently ongoing (see section 3.8) such that the mapping will become $R(s, t, \tilde{\phi})$ and $Z(s, t, \tilde{\phi})$, where $\tilde{\phi}$ does not have to be equivalent to ϕ, but can be a function of (R, Z, ϕ) (e.g., in the case of Boozer coordinates). The orthogonality assumption of ϕ with respect to s and t might be dropped in the context of this extension.

Appendix B. Normalization

- The growth rate $\gamma_{n,\mathrm{SI}}$ for harmonic n is calculated by $\ln[E_n(t_2)/E_n(t_1)]/[2(t_2 - t_1)]$ with E_n the mode energy and t_2 and t_1 the corresponding time points at which the energies are determined. The factor two appears in the denominator to obtain growth rates for amplitudes instead of energies and be comparable to linear codes.
- Heat diffusion coefficients in JOREK are defined absorbing the factor $\gamma - 1$, which would normally appear in the energy equation for diffusion terms. This means, JOREK diffusivities need to be multiplied by a factor 1.5 (at $\gamma = 5/3$) when comparing them to the usual definitions.
- Note that $\chi_{\mathrm{SI}}(\mathrm{m}^2\ \mathrm{s}^{-1}) = K_{\mathrm{SI}}/\rho_{\mathrm{SI}}$ where $\kappa_{\mathrm{SI}}(\mathrm{m}^{-1}\ \mathrm{s}^{-1}) = n_{\mathrm{SI}}\chi_{\mathrm{SI}}$
- Thin wall resistivity: $\eta_{\mathrm{wall,thin,SI}}(\Omega) = \eta_{\mathrm{wall,SI}}(\Omega\ \mathrm{m})/d_{\mathrm{wall}}(\mathrm{m})$ with wall thickness d_{wall}; e.g. ITER: $\eta_{\mathrm{wall,thin,SI}} = \eta_{\mathrm{wall,SI}}/d_{\mathrm{wall}} = 8 \times 10^{-7}\ \Omega\mathrm{m}/(6\mathrm{cm}) = 1.33 \times 10^{-5}\ \Omega$.

Appendix C. Time stepping scheme

The time-integration of a set of equations of the form

$$\frac{\partial \mathbf{A}(\mathbf{u})}{\partial t} = \mathbf{B}(\mathbf{u}, t) \tag{61}$$

can be discretized by the general form (refer to [299, 300])

Connection between SI and normalized units		Description of the quantity
R_{SI}(m)	$= R$	Major radius
Z_{SI}(m)	$= Z$	Vertical coordinate
$\mathbf{B}_{SI}$(T)	$= \mathbf{B}$	Magnetic field vector
$\mathbf{E}_{SI}$(V m^{-1})	$= \mathbf{E}/\sqrt{\mu_0 \rho_0}$	Electric field vector
Ψ_{SI}(T m^2)	$= \Psi$	Poloidal magnetic flux
$j_{\phi,SI}$(A m^{-2})	$= -j/(R\mu_0)$	Toroidal current density; $j_{\phi,SI} = \mathbf{j}_{SI} \cdot \hat{e}_\phi$
n_{SI}(m^{-3})	$= \rho n_0$	Particle density
ρ_{SI}(kg m^{-3})	$= \rho \rho_0$	Mass density = ion mass × particle density
T_{SI}(K)	$= T(k_B{}^{-1}\mu_0 n_0)$	Temperature = electron + ion temperature
T_{ev}(eV)	$= T/(e\mu_0 n_0)$	Temperature = electron + ion temperature
FF'_{SI} (T rad)	$= FF'$	Poloidal current stream function $F = RB_\phi$ and $' = \mathrm{d}/\mathrm{d}\psi$
p_{SI}(N m^{-2})	$= \rho T/\mu_0$	Plasma pressure
$\mathbf{v}_{SI}$(m s^{-1})	$= \mathbf{v}/\sqrt{\mu_0 \rho_0}$	Velocity vector
$v_{\parallel,SI}$(m s^{-1})	$= v_\parallel \cdot B_{SI}/\sqrt{\mu_0 \rho_0}$	Parallel velocity component, where $B_{SI} = |\mathbf{B}_{SI}|$
u_{SI}(m s^{-1})	$= u/\sqrt{\mu_0 \rho_0}$	Ru is the velocity stream function, $F_0 u$ is the potential
$\omega_{\phi,SI}$(m^{-1} s^{-1})	$= \omega/\sqrt{\mu_0 \rho_0}$	Toroidal vorticity
t_{SI}(s)	$= t \cdot \sqrt{\mu_0 \rho_0}$	Time
γ_{SI}(s^{-1})	$= \gamma/\sqrt{\mu_0 \rho_0}$	Growth rate of an MHD mode
η_{SI}(Ω m)	$= \eta \cdot \sqrt{\mu_0/\rho_0}$	Resistivity
ν_{SI}(kg m^{-1} s^{-1})	$= \nu \cdot \sqrt{\rho_0/\mu_0}$	Dynamic viscosity
$\tilde{\nu}_{SI}$ [m^2 s^{-1}]	$= \tilde{\nu}_{SI}/\rho_{SI}$	Kinematic viscosity (ρ_{SI} is the local value)
D_{SI}(m^2 s^{-1})	$= D/\sqrt{\mu_0 \rho_0}$	Particle diffusivity ($\parallel$ or $\perp$); usually, $D_\parallel = 0$
K_{SI}(kg m^{-1} s^{-1})	$= K \cdot \sqrt{\rho_0/\mu_0}/(\gamma - 1)$	Heat diffusivity ($\parallel$ or $\perp$)
$S_{T,SI}$(W m^{-3})	$= S_T/\sqrt{\mu_0^3 \rho_0}$	Heat source
$S_{\rho,SI}$(kg s^{-1} m^{-3})	$= S_\rho \cdot \sqrt{\rho_0/\mu_0}$	Particle source
$\eta_{wall,thin,SI}$(Ω)	$= \eta_{wall,thin} \cdot \sqrt{\mu_0/\rho_0}$	Resistivity of conducting structures
$R_{ion/rec,SI}$(m^{-3} s^{-1})	$= R_{ion/rec}/(\sqrt{\mu_0 \rho_0} n_0)$	Ionisation and recombination rate
$E_{ion,SI}$(J)	$= \xi_{ion}/(\frac{2}{3}\mu_0 n_0)$	Ionisation energy
$L_{rad,SI}$(W m^3)	$= L_{rad}/(\frac{2}{3}\sqrt{\rho_0/\mu_0} n_0^2 \frac{m_i}{m_{imp}})$	Radiation rate (impurity fluid model)
$P_{rad,SI}$(W m^{-3})	$= P_{rad}/(\frac{2}{3}\sqrt{\rho_0/\mu_0})$	Radiation power density (impurity fluid model)
q_{SI}(A s)	$= q\sqrt{\mu_0 \rho_0}$	Particle charge
$\mu_{neo,SI}$(s^{-1})	$= \mu_{neo}/\sqrt{\rho_0 \mu_0}$	Neoclassical friction rate

$$(1 + \xi)\mathbf{A}^{n+1} - (1 + 2\xi)\mathbf{A}^n + \xi\mathbf{A}^{n-1}$$
$$= \Delta t\left[\theta\mathbf{B}^{n+1} + (1 - \theta - \phi)\mathbf{B}^n - \phi\mathbf{B}^{n-1}\right], \quad (62)$$

which guarantees second-order accuracy, if $\phi + \theta - \xi = 1/2$. Superscripts like B^n indicate at which timestep the corresponding expression is evaluated. The linearization $\mathbf{H}^{n+1} \approx \mathbf{H}^n + \partial\mathbf{H}/\partial\mathbf{u}|^n \cdot \delta\mathbf{u}^n$, with $\mathbf{H} = \mathbf{A}$ or $\mathbf{H} = \mathbf{B}$, which is described in reference [300], allows to rewrite equation (62) in the following way, where $\phi = 0$ has been chosen:

$$(1 + \xi)\left[\mathbf{A}^n + \left(\frac{\partial\mathbf{A}}{\partial\mathbf{u}}\right)^n \delta\mathbf{u}^n\right] - (1 + \xi)\mathbf{A}^n - \xi\mathbf{A}^n + \xi\mathbf{A}^{n-1}$$
$$= \Delta t\left[\theta\left(\mathbf{B}^n + \left(\frac{\partial\mathbf{B}}{\partial\mathbf{u}}\right)^n \delta\mathbf{u}^n\right) + (1 - \theta)\mathbf{B}^n\right]. \quad (63)$$

Here, $\delta\mathbf{u}^n \equiv \mathbf{u}^{n+1} - \mathbf{u}^n$. After some simplifications, and using the backward linearization $\mathbf{H}^{n-1} \approx \mathbf{H}^n - \partial\mathbf{H}/\partial\mathbf{u}|^n \cdot \delta\mathbf{u}^{n-1}$, one obtains

$$\left[(1 + \xi)\left(\frac{\partial\mathbf{A}}{\partial\mathbf{u}}\right)^n - \Delta t\theta\left(\frac{\partial\mathbf{B}}{\partial\mathbf{u}}\right)^n\right]\delta\mathbf{u}^n$$
$$= \Delta t\mathbf{B}^n + \xi\left(\frac{\partial\mathbf{A}}{\partial\mathbf{u}}\right)^n \delta\mathbf{u}^{n-1}, \quad (64)$$

which is the time-integration scheme implemented in JOREK. Certain parameter choices correspond to well-known time integration methods: Crank–Nicolson is selected by $\theta = 1/2$ and $\xi = 0$, BDF2 (Gears) is selected by $\theta = 1$ and $\xi = 1/2$, and first order implicit Euler method (not used in production) corresponds to $\theta = 1$ and $\xi = 0$. The linearization shown above can also be replaced by Newton iterations with a beneficial effect onto non-linear stability in certain situation, as demonstrated in reference [64]. However, this is not implemented in the present code version.

Appendix D. Setup for testing anisotropic heat transport

For the tests shown in section 4.2, the following setup was used. The plasma cross section is circular with a major radius of $R_{axis} = 100$ m and a minor radius of 1 m. The magnetic configuration is initialized by the poloidal flux distribution $\Psi = 1 - [(R - R_{axis})^2 + Z^2]^2$ T m^{-2} and $F_0 = 200$ T m. The applied perturbation is given by $\tilde{\Psi} = 0.01$T$[Z \sin(\phi) + (R - R_{axis}) \cos(\phi)]$. The simulation domain is rectangular 2 m by 2 m for the non-aligned grid and circular with radius 1 m for the aligned grid. The temperature at the boundary of the

computational domain is fixed at zero via Dirichlet boundary conditions. For the following, (normalized) pre-factors are omitted since they do not affect the results due to the self-similarity of the solution and since only relative errors are discussed. To establish a steady state temperature distribution, a source $S_T = \{0.5 - 0.5 \tanh[(\Psi_n - 0.1)/0.1]\}$ is applied in the plasma centre. The perpendicular and parallel heat conduction coefficients are spatially constant and their ratio is varied for the tests. The density distribution is spatially constant. Convergence is tested by assessing the steady state value of the axis temperature that establishes.

ORCID iDs

M. Hoelzl https://orcid.org/0000-0001-7921-9176

V. Bandaru https://orcid.org/0000-0003-4096-1407

D. Bonfiglio https://orcid.org/0000-0003-2638-317X

A. Cathey https://orcid.org/0000-0001-7693-5556

S. Futatani https://orcid.org/0000-0001-5742-5454

D. Hu https://orcid.org/0000-0001-7099-5734

L. Kos https://orcid.org/0000-0002-1790-7093

J.A. Morales https://orcid.org/0000-0002-3140-0504

K. Särkimäki https://orcid.org/0000-0001-9156-2559

S.F. Smith https://orcid.org/0000-0003-2319-0356

References

[1] Huysmans G.T.A. and Czarny O. 2007 MHD stability in X-point geometry: simulation of ELMs *Nucl. Fusion* **47** 659

[2] Czarny O. and Huysmans G. 2008 Bezier surfaces and finite elements for MHD simulations *J. Comput. Phys.* **227** 7423–45

[3] Huijsmans G.T.A., Chang C.S., Ferraro N., Sugiyama L., Waelbroeck F., Xu X.Q., Loarte A. and Futatani S. 2015 Modelling of edge localised modes and edge localised mode control *Phys. Plasmas* **22** 021805

[4] Pamela S.J.P. *et al* (J ET Contributors, The MAST Team, The AUG Team, The JOREK Team) 2018 ELM and ELM-control simulations *Proc. of the 27th IAEA Fusion Energy Conf.* (Gandhinagar, India) p OV/4–4 (https://nucleus.iaea.org/sites/fusionportal/Shared%20Documents/FEC%202018/fec2018-preprints/preprint0579.pdf)

[5] Hoelzl M. *et al* (ASDEX Upgrade Team, EUROfusion MST1 Team) 2018 Insights into type-I edge localized modes and edge localized mode control from JOREK non-linear magneto-hydrodynamic simulations *Contrib. Plasma Phys.* **58** 518

[6] Hoelzl M. *et al* 2018 Simulating tokamak edge instabilities: advances and challenges *45th European Physical Society Conf. on Plasma Physics (EPS)* (Prague, Czech Republic) Invited Talk I5.J601 http://2.ipp.mpg.de/&tnqx7e;mhoelzl/2018-07-EPS-mhoelzl-slides.pdf)

[7] Dudson B.D., Umansky M.V., Xu X.Q., Snyder P.B. and Wilson H.R. 2009 BOUT++: a framework for parallel plasma fluid simulations *Comput. Phys. Commun.* **180** 1467–80

[8] Huysmans G.T.A., Pamela S., van der Plas E. and Ramet P. 2009 Non-linear MHD simulations of edge localized modes (ELMs) *Plasma Phys. Control. Fusion* **51** 124012

[9] Todo Y. and Sato T. 1998 Linear and nonlinear particle-magnetohydrodynamic simulations of the toroidal Alfvén eigenmode *Phys. Plasmas* **5** 1321–7

[10] Park W., Belova E.V., Fu G.Y., Tang X.Z., Strauss H.R. and Sugiyama L.E. 1999 Plasma simulation studies using multilevel physics models *Phys. Plasmas* **6** 1796–803

[11] Jardin S.C., Breslau J. and Ferraro N. 2007 A high-order implicit finite element method for integrating the two-fluid magnetohydrodynamic equations in two dimensions *J. Comput. Phys.* **226** 2146–74

[12] Jardin S.C., Ferraro N., Breslau J. and Chen J. 2012 Multiple timescale calculations of sawteeth and other global macroscopic dynamics of tokamak plasmas *Comput. Sci. Discovery* **5** 014002

[13] Ferraro N.M., Jardin S.C., Lao L.L., Shephard M.S. and Zhang F. 2016 Multi-region approach to free-boundary three-dimensional tokamak equilibria and resistive wall instabilities *Phys. Plasmas* **23** 056114

[14] Glasser A.H., Sovinec C.R., Nebel R.A., Gianakon T.A., Plimpton S.J., Chu M.S. and Schnack D.D. (The NIMROD Team) 1999 The NIMROD code: a new approach to numerical plasma physics *Plasma Phys. Control. Fusion* **41** A747–55

[15] Sovinec C.R. *et al* (The NIMROD Team) 2004 Nonlinear magnetohydrodynamics simulation using high-order finite elements *J. Comput. Phys.* **195** 355–86

[16] Lütjens H. and Luciani J.-F. 2008 The XTOR code for nonlinear 3D simulations of MHD instabilities in tokamak plasmas *J. Comput. Phys.* **227** 6944–66

[17] Lütjens H. and Luciani J.-F. 2010 XTOR-2F: a fully implicit Newton–Krylov solver applied to nonlinear 3D extended MHD in tokamaks *J. Comput. Phys.* **229** 8130–43

[18] Hans Goedbloed J.P. and Poedts S. 2004 *Principles of Magnetohydrodynamics: With Applications to Laboratory and Astrophysical Plasmas* (Cambridge: Cambridge University Press)

[19] Goedbloed J.P., Keppens R. and Poedts S. 2010 *Advanced Magnetohydrodynamics: With Applications to Laboratory and Astrophysical Plasmas* (Cambridge: Cambridge University Press)

[20] Schnack D.D. 2009 *Lectures in Magnetohydrodynamics* (Berlin: Springer)

[21] Biskamp D. 1993 *Nonlinear Magnetohydrodynamics* (Cambridge: Cambridge University Press)

[22] Freidberg J.P. 2014 *Ideal MHD* (Cambridge: Cambridge University Press)

[23] Zohm H. 2014 *Magnetohydrodynamic Stability of Tokamaks* (New York: Wiley)

[24] Stephen J. 2010 *Computational Methods in Plasma Physics* (Boca Raton, FL: CRC Press)

[25] Alfvén H. 1942 Existence of electromagnetic-hydrodynamic waves *Nature* **150** 405

[26] Gohil P. *et al* 1988 Study of giant edge-localized modes in DIII-D and comparison with ballooning theory *Phys. Rev. Lett.* **61** 1603–6

[27] Huysmans G.T.A. 1995 Influence of edge currents and pressure gradients on MHD stability of low-*n* external kink modes *22nd European Physical Society Conf. on Plasma Physics (EPS)* (Bournemouth, UK) p I.201 (http://euro-fusionscipub.org/wp-content/uploads/2014/11/JETP95037.pdf. Pages 33ff. in the PDF document)

[28] Connor J.W. 1998 Edge-localized modes—physics and theory *Plasma Phys. Control. Fusion* **40** 531–42

[29] Snyder P.B. *et al* 2004 ELMs and constraints on the H-mode pedestal: peeling-ballooning stability calculation and comparison with experiment *Nucl. Fusion* **44** 320

[30] Abdelhamid H.M., Kawazura Y. and Yoshida Z. 2015 Hamiltonian formalism of extended magnetohydrodynamics *J. Phys. A: Math. Theor.* **48** 235502

[31] Spitzer L. and Härm R. 1953 Transport phenomena in a completely ionized gas *Phys. Rev.* **89** 977–81

[32] Malone R.C., McCrory R.L. and Morse R.L. 1975 Indications of strongly flux-limited electron thermal conduction in laser–target experiments *Phys. Rev. Lett.* **34** 721–4

[33] Sauter O., Angioni C. and Lin-Liu Y.R. 1999 Neoclassical conductivity and bootstrap current formulas for general axisymmetric equilibria and arbitrary collisionality regime *Phys. Plasmas* **6** 2834–9

[34] Strauss H.R. 1983 Finite-aspect-ratio MHD equations for tokamaks *Nucl. Fusion* **23** 649–55

[35] Strauss H.R. 1997 Reduced MHD in nearly potential magnetic fields *J. Plasma Phys.* **57** 83–7

[36] Huysmans G.T.A. 2005 External kink (peeling) modes in X-point geometry *Plasma Phys. Control. Fusion* **47** 2107–21

[37] Nardon E., Bécoulet M., Huysmans G. and Czarny O. 2007 Magnetohydrodynamics modelling of H-mode plasma response to external resonant magnetic perturbations *Phys. Plasmas* **14** 092501

[38] Oden J.T., Duarte C.A.M. and Zienkiewicz O.C. 1998 A new cloud-based hp finite element method *Comput. Methods Appl. Mech. Eng.* **153** 117–26

[39] Huysmans G.T.A., Goedbloed J.P. and Kerner W. 1991 Isoparametric bicubic Hermite elements for solution of the Grad–Shafranov equation *Proc. CP90—Conf. on Computational Physics* (1-13 September 1990 Amsterdam) (Singapor: World Scientific) pp 371–6

[40] Huysmans G.T.A. and Czarny O. 2006 Implementation of cubic Bezier elements in the non-linear MHD code JOREK (technical report, project ASTER) *Technical Report* CEA (http://aster.gforge.inria.fr/reports/huysmans_report_aster_001.pdf)

[41] Huysmans G.T.A. 2006 Implementation of an iterative solver in the non-linear MHD code JOREK *Technical Report* ANR (http://aster.gforge.inria.fr/reports/huysmans_report_aster_002.pdf)

[42] JOREK 2014–2021 JOREK website https://jorek.eu (accessed: 1 March 2021)

[43] Sauter O. and Medvedev S.Y. 2013 Tokamak coordinate conventions *Comput. Phys. Commun.* **184** 293–302

[44] Hazeltine R.D. and Meiss J.D. 2013 *Plasma Confinement* (New York: Dover)

[45] Pamela S.J.P. *et al* 2017 Recent progress in the quantitative validation of JOREK simulations of ELMs in JET *Nucl. Fusion* **57** 076006

[46] Callen J.D. 2009 Viscous forces due to collisional parallel stresses for extended MHD codes *Technical Report UW-CPTC 09-6R* University of Wisconsin (http://homepages.cae.wisc.edu/&tnqx7e;callen/UW-CPTC_09-6R.pdf)

[47] Goldberger M.L., Chew G.F. and Low F.E. 1956 The Boltzmann equation and the one-fluid hydromagnetic equations in the absence of particle collisions *Proc. R. Soc.* A **236** 112

[48] Gianakon T.A., Kruger S.E. and Hegna C.C. 2002 Heuristic closures for numerical simulations of neoclassical tearing modes *Phys. Plasmas* **9** 536–47

[49] Orain F. *et al* 2013 Non-linear magnetohydrodynamic modeling of plasma response to resonant magnetic perturbations *Phys. Plasmas* **20** 102510

[50] Orain F. 2014 Edge localized mode control by resonant magnetic perturbations in tokamak plasmas *PhD Thesis* Physique et Sciences de la Matière(http://theses.fr/2014AIXM4749/document)

[51] Hazeltine R.D., Kotschenreuther M. and Morrison P.J. 1985 A four-field model for tokamak plasma dynamics *Phys. Fluids* **28** 2466–77

[52] Schnack D.D., Cheng J., Barnes D.C. and Parker S.E. 2013 Comparison of kinetic and extended magnetohydrodynamics computational models for the linear ion temperature gradient instability in slab geometry *Phys. Plasmas* **20** 062106

[53] Guillard H. 2015 The mathematical theory of reduced MHD models for fusion plasmas *Research Report RR-8715, Inria* (https://hal.inria.fr/hal-01145009)

[54] Kruger S.E., Hegna C.C. and Callen J.D. 1998 Generalized reduced magnetohydrodynamic equations *Phys. Plasmas* **5** 4169–82

[55] Greene J.M. and Johnson J.L. 1961 Determination of hydromagnetic equilibria *Phys. Fluids* **4** 875–90

[56] Kadomtsev B.B. and Pogutse O.P. 1974 Nonlinear helical perturbations of a plasma in the tokamak *Sov. Phys. - JETP* **38** 283–90

[57] Strauss H.R. 1976 Nonlinear, three-dimensional magnetohydrodynamics of noncircular tokamaks *Phys. Fluids* **19** 134–40

[58] Strauss H.R. 1977 Dynamics of high β tokamaks *Phys. Fluids* **20** 1354–60

[59] Strauss H.R. 1980 Stellarator equations of motion *Plasma Phys.* **22** 733–45

[60] Oughton S., Matthaeus W.H. and Dmitruk P. 2017 Reduced MHD in astrophysical applications: two-dimensional or three-dimensional? *Astrophys. J.* **839** 2

[61] Park W., Monticello D.A., White R.B. and Jardin S.C. 1980 Non-linear saturation of the internal kink mode *Nucl. Fusion* **20** 1181–5

[62] Izzo R., Monticello D.A., DeLucia J., Park W. and Ryu C.M. 1985 Reduced equations for finite beta tearing modes in tokamaks *Phys. Fluids* **28** 903–11

[63] Breslau J., Ferraro N. and Jardin S. 2009 Some properties of the M3D-C1 form of the three-dimensional magnetohydrodynamics equations *Phys. Plasmas* **16** 092503

[64] Franck E., Hölzl M., Lessig A. and Sonnendrücker E. 2015 Energy conservation and numerical stability for the reduced MHD models of the non-linear JOREK code *ESAIM: Math. Modelling Numer. Anal.* **49** 1331–65

[65] Pamela S.J.P., Bhole A., Huijsmans G.T.A., Nkonga B., Hoelzl M., Krebs I. and Strumberger E. 2020 Extended full-MHD simulation of non-linear instabilities in tokamak plasmas *Phys. Plasmas* **27** 102510

[66] Stangeby P.C. 2000 *The Plasma Boundary of Magnetic Fusion Devices* (Boca Raton, FL: CRC Press)

[67] Artola F.J. *et al* (JOREK Team, COMPASS Team) 2020 Simulations of COMPASS vertical displacement events with a self-consistent model for halo currents including neutrals and sheath boundary conditions *Proc. of the Joint Varenna—Lausanne International Workshop on Theory of Fusion Plasmas* (Lausanne, Switzerland 12–16 October 2020)

[68] Graves J.P., Zullino D., Brunetti D., Lanthaler S. and Wahlberg C. 2019 Reduced models for parallel magnetic field fluctuations and their impact on pressure gradient driven MHD instabilities in axisymmetric toroidal plasmas *Plasma Phys. Control. Fusion* **61** 104003

[69] Zeiler A. 1988 Tokamak edge turbulence *Technical Report IPP Report 5/88* IPP Garching (http://hdl.handle.net/11858/00-001M-0000-0027-602C-C)

[70] Nikulsin N., Hoelzl M., Zocco A., Lackner K. and Guenter S. 2021 Testing of the new JOREK stellarator-capable model in the tokamak limit *J. Plasma Phys.* (in press)

[71] Nikulsin N., Hoelzl M., Zocco A., Lackner K. and Günter S. 2019 A three-dimensional reduced MHD model consistent with full MHD *Phys. Plasmas* **26** 102109

[72] Mitterauer V., Schwarz N., Artola F.J. and Hoelzl M. 2021 Implementation and verification for 3D Coils in JOREK–STARWALL (private communication)

[73] Pamela S. 2010 Simulation magnéto-hydro-dynamiques des edge-localised-modes dans un tokamak *PhD Thesis* Université de Provence (https://inis.iaea.org/collection/NCLCollectionStore/_Public/43/127/43127242.pdf)

[74] Reux C., Huysmans G., Bucalossi J. and Becoulet M. 2011 Non-linear simulations of disruption mitigation using massive gas injection on tore supra *38th European Physical Society Conf. on Plasma Physics (EPS)* (Strasbourg, France) p O3.117 (http://ocs.ciemat.es/EPS2011PAP/pdf/O3.117.pdf)

[75] Fil A. 2015 Modélisation des disruptions déclenchées par injection massive de gaz dans les plasmas de tokamaks *PhD Thesis* Aix-Marseille University (http://theses.fr/2015AIXM4040)

[76] Smith S.F., Pamela S.J.P., Hoelzl M., Huijsmans G.T.A., Kirk A., Moulton D., Thornton A.J. and Wilson H.R. 2019 Burn-through due to plasma eruptions in the mast-u Super-X divertor *46th European Physical Society Conf. on Plasma Physics (EPS)* (Milan, Italy) p P2.1037 (http://ocs.ciemat.es/EPS2019PAP/pdf/P2.1037.pdf)

[77] Huijsmans G.T.A., van Vugt D.C., Franssen S., Korving S.Q. and Becoulet M. 2019 Non-linear MHD simulations of ELMs in a high recycling devertor *46th European Physical Society Conf. on Plasma Physics (EPS)* (Milan, Italy) p P2.1059 (https://hal-cea.archives-ouvertes.fr/cea-02555093/document)

[78] Nardon E., Fil A., Hoelzl M. and Huijsmans G. (JET Contributors) 2017 Progress in understanding disruptions triggered by massive gas injection via 3D non-linear MHD modelling with JOREK *Plasma Phys. Control. Fusion* **59** 014006

[79] Hu D., Nardon E., Huijsmans G.T.A., Lehnen M. and van Vugt D.C. (JET Contributors) 2018 JOREK simulations of shattered pellet injection with high *Z* impurities *45th European Physical Society Conf. on Plasma Physics (EPS)* (Prague, Czech Republic) p P4.1043 (http://ocs.ciemat.es/EPS2018PAP/pdf/P4.1043.pdf)

[80] Hu D. *et al* 2021 Radiation asymmetry and MHD destabilization during the thermal quench after impurity shattered pellet injection *Nucl. Fusion* **61** 026015

[81] Lyons B.C., Kim C.C., Liu Y.Q., Ferraro N.M., Jardin S.C., McClenaghan J., Parks P.B. and Lao L.L. 2019 Axisymmetric benchmarks of impurity dynamics in extended-magnetohydrodynamic simulations *Plasma Phys. Control. Fusion* **61** 064001

[82] Parks P.B. and Turnbull R.J. 1978 Effect of transonic flow in the ablation cloud on the lifetime of a solid hydrogen pellet in a plasma *Phys. Fluids* **21** 1735–41

[83] Sergeev V.Y., Bakhareva O.A., Kuteev B.V. and Tendler M. 2006 Studies of the impurity pellet ablation in the high-temperature plasma of magnetic confinement devices *Plasma Phys. Rep.* **32** 363–77

[84] Parks P.B., Bosviel N. and Samulyak 2020 Numerical studies of neon pellet ablation in tokamaks *46th European Physical Society Conf. on Plasma Physics (EPS)* (Milan, Italy) p P1.1102 (http://ocs.ciemat.es/EPS2019PAP/pdf/P1.1102.pdf)

[85] Parks P.B. 2019 The ablation rate of light-element pellets with a kinetic treatment for penetration of plasma electrons through the ablation cloud (private communication)

[86] Futatani S. *et al* 2014 Non-linear MHD modelling of ELM triggering by pellet injection in DIII-D and implications for ITER *Nucl. Fusion* **54** 073008

[87] Hu D., Nardon E., Lehnen M., Huijsmans G.T.A. and van Vugt D.C. 2018 3D non-linear MHD simulation of the MHD response and density increase as a result of shattered pellet injection *Nucl. Fusion* **58** 126025

[88] Hoelzl M., Merkel P., Huysmans G.T.A., Nardon E., Strumberger E., McAdams R., Chapman I., Günter S. and Lackner K. 2012 Coupling JOREK and STARWALL codes for non-linear resistive-wall simulations *J. Phys.: Conf. Ser.* **401** 012010

[89] Merkel P. and Strumberger E. 2015 Linear MHD stability studies with the STARWALL code (arXiv:1508.04911)

[90] Heumann H. and Rapetti F. 2017 A finite element method with overlapping meshes for free-boundary axisymmetric plasma equilibria in realistic geometries *J. Comput. Phys.* **334** 522–40

[91] Hoelzl M., Schwarz N., Artola F.J. and Mitterauer V. 2017 Implementation of self-consistent coils in JOREK-STARWALL (private communication)

[92] Artola Such F.J. 2018 Free-boundary simulations of MHD plasma instabilities in tokamaks *PhD Thesis* Université Aix Marseille (https://hal-amu.archives-ouvertes.fr/tel-02012234v1)

[93] Zakharov L.E., Atanasiu C.V., Lackner K., Hoelzl M. and Strumberger E. 2015 Electromagnetic thin-wall model for simulations of plasma wall-touching kink and vertical modes *J. Plasma Phys.* **81** 515810610

[94] Atanasiu C.V., Zakharov L.E., Lackner K., Hoelzl M., Strumberger E. and Nicolici S. 2016 Simulation of surface currents excited by plasma wall-touching kink and vertical modes in tokamak *2016 Int. Symp. on Fundamentals of Electrical Engineering (ISFEE)* 30 June-2 July 2016 pp 1–4

[95] Atanasiu Č.V., Zakharov L.E., Lackner K. and Hoelzl M. 2018 Simulation of the electromagnetic wall response during vertical displacement events (VDE) in ITER tokamak *J. Phys.: Conf. Ser.* **1141** 012065

[96] Krebs I., Artola F.J., Sovinec C.R., Jardin S.C., Bunkers K.J., Hoelzl M. and Ferraro N.M. 2020 Axisymmetric simulations of vertical displacement events in tokamaks: a benchmark of M3D-C1, NIMROD, and JOREK *Phys. Plasmas* **27** 022505

[97] Atanasiu C.V., Zakharov L.E., Hoelzl M. and Gerasimov S. 2021 A mathematical model for calculation of the influence of ferromagnetic components in vertical displacement events and stability simulations of tokamak plasmas *J. Phys.: Conf. Ser.* **1730** 012115

[98] Artola F.J., Sovinec C.R., Jardin S.C., Hoelzl M., Krebs I. and Clauser C. 2021 3D simulations of vertical displacement events in tokamaks: a benchmark of M3D-C1, NIMROD and JOREK *Phys. Plasmas* (in press)

[99] Mochalskyy S., Hoelzl M. and Hatzky R. 2017 MPI parallelization of the resistive wall code STARWALL: report of the EUROfusion high level support team project JORSTAR (arXiv:1609.07441)

[100] Delzanno G.L. and Camporeale E. 2013 On particle movers in cylindrical geometry for particle-in-cell simulations *J. Comput. Phys.* **253** 259–77

[101] Sommariva C., Nardon E., Beyer P., Hoelzl M., Huijsmans G.T.A. and van Vugt D. 2017 Test particles dynamics in the JOREK 3D non-linear MHD code and application to electron transport in a disruption simulation *Nucl. Fusion* **58** 016043

[102] Sommariva C. 2017 Dynamique des particules tests dans des simulations de magnétohydrodynamiquenon-linéaire 3D et application à la formation des électrons découplés pendant les disruptions plasma *PhD Thesis* Aix Marseille Université (https://hal.archives-ouvertes.fr/tel-01813986)

[103] Takizuka T. and Abe H. 1977 A binary collision model for plasma simulation with a particle code *J. Comput. Phys.* **25** 205–19

[104] Homma Y. and Hatayama A. 2011 Numerical modeling of thermal force in a plasma for test-ion transport simulation based on Monte Carlo binary collision model *J. Comput. Phys.* **231** 3211–27

[105] Homma Y. and Hatayama A. 2013 Numerical modeling of the thermal force in a plasma for test-ion transport simulation based on a Monte Carlo binary collision model (II)—thermal forces due to temperature gradients parallel

and perpendicular to the magnetic field *J. Comput. Phys.* **250** 206–23

[106] Eckstein W. 2002 Calculated sputtering, reflection and range values *Technical Report IPP Report 9/132* IPP Garching (https://pure.mpg.de/rest/items/item_2138250/component/file_2138249/content)

[107] Eckstein W. 2008 Sputtering yields *Vacuum* **82** 930–4

[108] Eckstein W. 2011 Sputtered energy coefficient and sputtering yeild *Technical Report IPP Report 17/29* IPP Garching (https://pure.mpg.de/rest/items/item_2140248/component/file_2140247/content)

[109] Kallenbach A. *et al* 2005 Tokamak operation with high-Z plasma facing components *Plasma Phys. Control. Fusion* **47** B207–22

[110] Bandaru V., Hoelzl M., Artola F.J., Papp G. and Huijsmans G.T.A. 2019 Simulating the nonlinear interaction of relativistic electrons and tokamak plasma instabilities: implementation and validation of a fluid model *Phys. Rev. E* **99** 063317

[111] Bandaru V., Hoelzl M. *et al* 2020 Including the re fluid in JOREK into the MHD force balance equation (private communication)

[112] Haverkort J.W. 2013 Magnetohydrodynamic waves and instabilities in rotating tokamak plasmas *PhD Thesis* Department of Mathematics and Computer Science, Technische Universiteit Eindhoven (https://pure.tue.nl/ws/portalfiles/portal/3447608/750676.pdf)

[113] Haverkort J.W., de Blank H.J., Huysmans G.T.A., Pratt J. and Koren B. 2016 Implementation of the full viscoresistive magnetohydrodynamic equations in a nonlinear finite element code *J. Comput. Phys.* **316** 281–302

[114] Huijsmans G.T.A. and Nkonga B. 2018 Implementation of sheath boundary conditions and numerical stabilization methods for the JOREK full MHD model (private communication)

[115] Garbet X. *et al* 2001 Global simulations of ion turbulence with magnetic shear reversal *Phys. Plasmas* **8** 2793–803

[116] Zielinski J., Smolyakov A.I., Beyer P., Becoulet M., Huijsmans G., Garbet X. and Benkadda S. 2017 On fluid modeling of ion temperature gradient (ITG) driven modes *44th European Physical Society Conf. on Plasma Physics (EPS)* (Belfast, Northern Ireland) p P5.163 (http://ocs.ciemat.es/EPS2017PAP/pdf/P5.163.pdf)

[117] Zielinski J., Becoulet M., Smolyakov A.I., Garbet X., Huijsmans G.T.A., Beyer P. and Benkadda S. 2020 Global ITG eigenmodes: from ballooning angle and radial shift to Reynolds stress and nonlinear saturation *Phys. Plasmas* **27** 072507

[118] Dimits A.M. *et al* 2000 Comparisons and physics basis of tokamak transport models and turbulence simulations *Phys. Plasmas* **7** 969–83

[119] Merlo G., Dominski J., Bhattacharjee A., Chang C.S., Jenko F., Ku S., Lanti E. and Parker S. 2018 Cross-verification of the global gyrokinetic codes GENE and XGC *Phys. Plasmas* **25** 062308

[120] Biancalani A. *et al* 2017 Cross-code gyrokinetic verification and benchmark on the linear collisionless dynamics of the geodesic acoustic mode *Phys. Plasmas* **24** 062512

[121] Lanthaler S., Graves J.P., Pfefferlé D. and Cooper W.A. 2019 Guiding-centre theory for kinetic-magnetohydrodynamic modes in strongly flowing plasmas *Plasma Phys. Control. Fusion* **61** 074006

[122] Lu Z.X., Lauber P., Hayward-Schneider T., Bottino A. and Hoelzl M. 2019 Development and testing of an unstructured mesh method for whole plasma gyrokinetic simulations in realistic tokamak geometry *Phys. Plasmas* **26** 122503

[123] Gruca M. 2020 Simulations of edge instabilities in TCV (private communication)

[124] Pamela S., Huijsmans G., Thornton A.J., Kirk A., Smith S.F., Hoelzl M. and Eich T. 2019 A wall-aligned grid generator for non-linear simulations of MHD instabilities in tokamak plasmas *Comput. Phys. Commun.* **243** 41–50

[125] Hughes T.J.R., Scovazzi G. and Franca L.P. 2004 *Multiscale and Stabilized Methods* (New York: Wiley)

[126] Billaud M., Gallice G. and Nkonga B. 2011 A simple stabilized finite element method for solving two phase compressible–incompressible interface flows *Comput. Methods Appl. Mech. Eng.* **200** 1272–90

[127] Jean D. 1984 A Taylor–Galerkin method for convective transport problems *Int. J. Numer. Methods Eng.* **20** 101–19

[128] Roig B. 2007 One-step Taylor–Galerkin methods for convection–diffusion problems *J. Comput. Appl. Math.* **204** 95–101

[129] Nkonga B., Tarcisio-Costa J. and Vides J. 2016 VMS finite element for MHD and reduced-MHD in tokamak plasmas *Technical Report, Inria Research Report* (https://hal.inria.fr/hal-01294788/)

[130] Hammett G.W. and Peterson J.L. 2008 Comparison of (some) algorithms for edge gyrokinetics *Gyrokinetic Turbulence Workshop* (Wolfgang Pauli Institute 15–19 September 2008) (https://thphys.physics.ox.ac.uk/people/AlexanderSchekochihin/wpi/workshop3_pdfs/hammett2.pdf)

[131] Fehér T.B., Hoelzl M., Latu G. and Huijsmans G.T.A. 2018 Performance analysis and optimization of the JOREK code for many-core CPUs (arXiv:1810.04413)

[132] Holod I., Hoelzl M., Verma P.S., Nies R. and Huijsmans G.T.A. (JOREK Team) 2021 New developments regarding the JOREK solver and physics based preconditioner *Plasma Phys. Control. Fusion* (private communication)

[133] Hénon P., Ramet P. and Roman J. 2002 PaStiX: a high-performance parallel direct solver for sparse symmetric positive definite systems *Parallel Comput.* **28** 301–21

[134] Amestoy P.R., Duff I.S., L'Excellent J.-Y. and Koster J. 2001 A fully asynchronous multifrontal solver using distributed dynamic scheduling *SIAM J. Matrix Anal. Appl.* **23** 15–41

[135] Amestoy P.R., Guermouche A., L'Excellent J.-Y. and Pralet S. 2006 Hybrid scheduling for the parallel solution of linear systems *Parallel Comput.* **32** 136–56

[136] Gupta A. 2002 Recent advances in direct methods for solving unsymmetric sparse systems of linear equations *ACM Trans. Math. Softw.* **28** 301–24

[137] Chevalier C. and Pellegrini F. 2008 PT-Scotch: a tool for efficient parallel graph ordering *Parallel Comput.* **34** 318–31

[138] Karypis G., Kirk S. and Kumar V. 2003 ParMETIS graph partitioning and sparse matrix ordering library *Technical Report* University of Minnesota Minneapolis (https://dev.ece.ubc.ca/projects/gpgpu-sim/export/96f6ad00d6d3e9a58b1d51edaac76d061c02fa82/ispass2009-benchmarks/DG/3rdParty/ParMetis-3.1/Manual/manual.pdf)

[139] Ghysels P., Xiaoye S.L., Gorman C. and Rouet F. 2017 A robust parallel preconditioner for indefinite systems using hierarchical matrices and randomized sampling *2017 IEEE Int. Parallel and Distributed Processing Symp. (IPDPS)* 29 May-2 June 2017 pp 897–906

[140] Nies R. and Hoelzl M. 2019 Testing performance with and without block low rank compression in MUMPS and the new PaStiX 6.0 for JOREK nonlinear MHD simulations (arXiv:1907.13442)

[141] Frayssé V., Giraud L. and Gratton S. 1997 A set of GMRES routines for real and complex arithmetics *Technical Report TR/PA/97/49* CERFACS Toulouse Cedex (https://pdfs.semanticscholar.org/bf99/0f2f085410ef17e851cd9bd4395a07bf7f5a.pdf)

[142] Latu G. *et al* 2012 Non regression testing for the JOREK code (arXiv:1211.3823)

[143] Hindenlang F., Maj O., Strumberger E., Rampp M. and Sonnendrücker E. 2019 GVEC: a newly developed 3D ideal MHD Galerkin variational equilibrium code *Simons Collaboration on Hidden Symmetries and Fusion Energy* (https://hiddensymmetries.princeton.edu/meetings/simons-hour-video-talk)

[144] Imbeaux F. *et al* 2015 Design and first applications of the ITER integrated modelling & analysis suite *Nucl. Fusion* **55** 123006

[145] Penko D., Kos L., Kos L., Bonnin X. and Pinches S.D. 2020 Post-processing for ITER scrape-off layer plasma simulations (SOLPS-ITER) in IMAS framework *J. Fusion Energy* **39** 202

[146] Penko D., Kos L., Huijsmans G.T.A. and Pinches S.D. 2019 The initial step towards JOREK integration in IMAS *28th Int. Conf. Nuclear Energy for New Europe* (Portorož, Slovenia 9–12 September 2019) pp 707.1–8

[147] Pamela S.J.P., Huijsmans G.T.A. and Hoelzl M. (The JOREK Team) 2021 A generalised formulation of *G*-continuous Bezier elements applied to non-linear MHD simulations *J. Comput. Phys.* (submitted)

[148] Strang G. and Fix G. 2008 *An Analysis of the Finite Element Method* 2nd edn (Cambridge: Cambridge University Press)

[149] Pamela S., Huysmans G. and Benkadda S. 2010 Influence of poloidal equilibrium rotation in MHD simulations of edge-localized modes *Plasma Phys. Control. Fusion* **52** 075006

[150] Günter S., Lackner K. and Tichmann C. 2007 Finite element and higher order difference formulations for modelling heat transport in magnetised plasmas *J. Comput. Phys.* **226** 2306–16

[151] Mikhailovskii A.B., Huysmans G.T.A., Kerner W.O.K. and Sharapov S.E. 1997 Optimization of computational mhd normal-mode analysis for tokamaks *Plasma Phys. Rep.* **23** 844

[152] Huysmans G.T.A., Sharapov S.E., Mikhailovskii A.B. and Kerner W. 2001 Modeling of diamagnetic stabilization of ideal magnetohydrodynamic instabilities associated with the transport barrier *Phys. Plasmas* **8** 4292–305

[153] Huysmans G.T.A., Goedbloed J.P. and Kerner W. 1993 Free boundary resistive modes in tokamaks *Phys. Fluids* B **5** 1545–58

[154] Strumberger E. and Günter S. 2017 CASTOR3D: linear stability studies for 2D and 3D tokamak equilibria *Nucl. Fusion* **57** 016032

[155] Pratt J., Huijsmans G.T.A. and Westerhof E. 2016 Early evolution of electron cyclotron driven current during suppression of tearing modes in a circular tokamak *Phys. Plasmas* **23** 102507

[156] Cappello S. and Biskamp D. 1996 Reconnection processes and scaling laws in reversed field pinch magnetohydrodynamics *Nucl. Fusion* **36** 571

[157] Bonfiglio D., Chacón L. and Cappello S. 2010 Nonlinear three-dimensional verification of the SPECYL and PIXIE3D magnetohydrodynamics codes for fusion plasmas *Phys. Plasmas* **17** 082501

[158] Dudson B.D., Xu X.Q., Umansky M.V., Wilson H.R. and Snyder P.B. 2011 Simulation of edge localized modes using BOUT++ *Plasma Phys. Control. Fusion* **53** 054005

[159] Smith S. 2019 Numerical simulations of edge localised mode instabilities in the MAST-U Super-X tokamak *PhD Thesis* University of York (http://etheses.whiterose.ac.uk/26275/)

[160] Morales J.A. *et al* 2016 Edge localized mode rotation and the nonlinear dynamics of filaments *Phys. Plasmas* **23** 042513

[161] Bécoulet M. *et al* 2017 Non-linear MHD modelling of edge localized modes dynamics in kstar *Nucl. Fusion* **57** 116059

[162] Bonnin X., Dekeyser W., Pitts R., Coster D., Voskoboynikov S. and Wiesen S. 2016 Presentation of the new SOLPS-ITER

code package for tokamak plasma edge modelling *Plasma Fusion Res.* **11** 1403102

[163] Smith S.F. *et al* 2020 Simulations of edge localised mode instabilities in MAST-U Super-X tokamak plasmas *Nucl. Fusion* **60** 066021

[164] Moulton D., Lipschultz B. and Harrison J. 2017 Detachment onset in MAST-U according to SOLPS-ITER *44th European Physical Society Conf. on Plasma Physics (EPS)* (Belfast, Northern Ireland) p O5.129 (http://ocs.ciemat.es/EPS2017PAP/pdf/O5.129.pdf)

[165] Papp G. *et al* (JET EFDA Contributors) 2013 The effect of ITER-like wall on runaway electron generation in JET *Nucl. Fusion* **53** 123017

[166] Hesslow L., Embréus O., Vallhagen O. and Fülöp T. 2019 Influence of massive material injection on avalanche runaway generation during tokamak disruptions *Nucl. Fusion* **59** 084004

[167] Matsuyama A., Aiba N. and Yagi M. 2017 Reduced fluid simulation of runaway electron generation in the presence of resistive kink modes *Nucl. Fusion* **57** 066038

[168] Hoelzl M. *et al* 2014 Non-linear simulations of MHD instabilities in tokamaks including eddy current effects and perspectives for the extension to halo currents *J. Phys.: Conf. Ser.* **561** 012011

[169] Schwarz N. 2020 Vertical displacement events in ASDEX upgrade *Technical Report* Max-Planck-Institut für Plasmaphysik

[170] Krebs I. *et al* 2018 Non-linear 3D simulations of vertical displacement events in tokamaks: progress on benchmark activities *Proc. of the 27th IAEA Fusion Energy Conf.* (Gandhinagar, India) p TH/P8–10 (https://nucleus.iaea.org/sites/fusionportal/Shared%20Documents/FEC%202018/fec2018-preprints/preprint0025.pdf)

[171] Artola F.J., Huijsmans G.T.A., Hoelzl M., Beyer P., Loarte A. and Gribov Y. 2018 Non-linear magnetohydrodynamic simulations of edge localised mode triggering via vertical position oscillations in ITER *Nucl. Fusion* **58** 096018

[172] van Vugt D.C. 2019 Nonlinear coupled MHD-kinetic particle simulations of heavy impurities in tokamak plasmas *PhD Thesis* Department of Applied Physics, Technische Universiteit Eindhoven (https://research.tue.nl/en/publications/nonlinear-coupled-mhd-kinetic-particle-simulations-of-heavy-impur)

[173] Tang W., Qin H. and Guan X. 2009 Variational symplectic algorithm for guiding center dynamics and its application in tokamak geometry *Phys. Plasmas* **16** 042510

[174] Kraus M. 2013 Variational integrators in plasma physics *PhD Thesis* Fakultat fur Mathematik, Technische Universitat Munich

[175] Särkimäki K., Embreus O., Nardon E. and Fülöp T. (JET Contributors) 2020 Assessing energy dependence of the transport of relativistic electrons in perturbed magnetic fields with orbit-following simulations *Nucl. Fusion* **60** 126050

[176] Varje J., Särkimäki K., Kontula J., Ollus P., Kurki-Suonio T., Snicker A., Hirvijoki E. and Äkäslompolo S. 2019 High-performance orbit-following code ASCOT5 for Monte Carlo simulations in fusion plasmas (unpublished) (arXiv:1908.02482)

[177] Zhang R., Liu J., Qin H., Wang Y., He Y. and Sun Y. 2015 Volume-preserving algorithm for secular relativistic dynamics of charged particles *Phys. Plasmas* **22** 044501

[178] Zohm H. 1996 Edge localized modes (ELMs) *Plasma Phys. Control. Fusion* **38** 105

[179] Loarte A. *et al* 2003 Characteristics of type I ELM energy and particle losses in existing devices and their extrapolation to ITER *Plasma Phys. Control. Fusion* **45** 1549

[180] Wenninger R. *et al* 2015 Advances in the physics basis for the European DEMO design *Nucl. Fusion* **55** 063003

[181] Xu X.Q., Xia T.Y., Yan N., Liu Z.X., Kong D.F., Diallo A., Groebner R.J., Hubbard A.E. and Hughes J.W. 2016 Toward integrated multi-scale pedestal simulations including edge-localized-mode dynamics, evolution of edge-localized-mode cycles, and continuous fluctuations *Phys. Plasmas* **23** 055901

[182] King J.R., Kruger S.E., Burrell K.H., Chen X., Garofalo A.M., Groebner R.J., Olofsson K.E.J., Pankin A.Y. and Snyder P.B. 2017 MHD modeling of a DIII-D low-torque QH-mode discharge and comparison to observations *Phys. Plasmas* **24** 055902

[183] Fil A., Kolemen E., Bortolon A., Ferraro N., Jardin S., Parks P.B., Lunsford R. and Maingi R. 2017 Modeling of lithium granule injection in NSTX with M3D-C1 *Nucl. Mater. Energy* **12** 1094–9

[184] Wilcox R.S., Wingen A., Cianciosa M.R., Ferraro N.M., Hirshman S.P., Paz-Soldan C., Seal S.K., Shafer M.W. and Unterberg E.A. 2017 Modeling of 3D magnetic equilibrium effects on edge turbulence stability during RMP ELM suppression in tokamaks *Nucl. Fusion* **57** 116003

[185] Moyer R.A. *et al* 2017 Validation of the model for ELM suppression with 3D magnetic fields using low torque ITER baseline scenario discharges in DIII-D *Phys. Plasmas* **24** 102501

[186] Canal G.P. *et al* 2017 M3D-C1 simulations of the plasma response to RMPs in NSTX-U single-null and snowflake divertor configurations *Nucl. Fusion* **57** 076007

[187] Wu Y.B., Xia T.Y., Zhong F.C., Zheng Z. and Liu J.B. 2018 Simulations of particle and heat fluxes in an ELMy h-mode discharge on EAST using BOUT++ code *Plasma Phys. Control. Fusion* **60** 055007

[188] Ebrahimi F. 2017 Nonlinear reconnecting edge localized modes in current-carrying plasmas *Phys. Plasmas* **24** 056119

[189] Xu G.S. *et al* 2019 Promising high-confinement regime for steady-state fusion *Phys. Rev. Lett.* **122** 255001

[190] Hu Q.M., Nazikian R., Grierson B.A., Logan N.C., Paz-Soldan C. and Yu Q. 2020 The role of edge resonant magnetic perturbations in edge-localized-mode suppression and density pump-out in low-collisionality DIII-D plasmas *Nucl. Fusion* **60** 076001

[191] Li M., Sun J., Wang Y. and Xia T. 2021 Nonlinear simulation of edge localized mode with pressure profile modified by pellet injection through a BOUT++ three-field mhd model *Nucl. Mater. Energy* **26** 100888

[192] Wagner F. *et al* 1982 Regime of improved confinement and high beta in neutral-beam-heated divertor discharges of the ASDEX tokamak *Phys. Rev. Lett.* **49** 1408–12

[193] Huijsmans G.T.A. and Loarte A. 2013 Non-linear MHD simulation of ELM energy deposition *Nucl. Fusion* **53** 123023

[194] Eich T., Sieglin B., Thornton A.J., Faitsch M., Kirk A., Herrmann A. and Suttrop W. 2017 ELM divertor peak energy fluence scaling to ITER with data from JET, MAST and ASDEX Upgrade *Nucl. Mater. Energy* **12** 84

[195] Ham C., Kirk A., Pamela S. and Wilson H. 2020 Filamentary plasma eruptions and their control on the route to fusion energy *Nat. Rev. Phys.* **2** 159–67

[196] Meyer H. *et al* 2019 Overview of physics studies on ASDEX Upgrade *Nucl. Fusion* **59** 112014

[197] Hoelzl M. and Günter S. (ASDEX Upgrade Team) 2011 Reduced MHD simulations of edge localized modes in ASDEX Upgrade *38th European Physical Society Conf. on Plasma Physics (EPS)* (Strasbourg, France) p P2.078 (http://ocs.ciemat.es/EPS2011PAP/pdf/P2.078.pdf)

[198] Hoelzl M., Guenter S., Wenninger R.P., Mueller W.-C., Huysmans G.T.A., Lackner K. and Krebs I. 2012 Reduced-magnetohydrodynamic simulations of toroidally and poloidally localized edge localized modes *Phys. Plasmas* **19** 082505

[199] Wenninger R.P. *et al* (The ASDEX Upgrade Team) 2012 Solitary magnetic perturbations at the ELM onset *Nucl. Fusion* **52** 114025

[200] Krebs I., Hölzl M., Lackner K. and Günter S. 2013 Nonlinear excitation of low-*n* harmonics in reduced magnetohydrodynamic simulations of edge-localized modes *Phys. Plasmas* **20** 082506

[201] Krebs I. 2012 Non-linear reduced MHD simulations of edge-localized modes in realistic ASDEX Upgrade geometry *Master's Thesis* Ludwig-Maximilians-Universität München (http://2.ipp.mpg.de/&tnqx7e;mhoelzl/ikrebs-thesis.pdf)

[202] Cathey A., Hoelzl M., Lackner K., Huijsmans G.T.A., Dunne M.G., Wolfrum E., Pamela S.J.P., Orain F. and Günter S. 2020 Non-linear extended MHD simulations of type-I edge localised mode cycles in ASDEX upgrade and their underlying triggering mechanism *Nucl. Fusion* **60** 124007

[203] Orain F. *et al* 2016 Non-linear modeling of the edge localized modes and their control by resonant magnetic perturbations in ASDEX Upgrade *43rd European Physical Society Conf. on Plasma Physics (EPS)* (Leuven, Belgium) p P1.017 (http://ocs.ciemat.es/EPS2016PAP/pdf/P1.017.pdf)

[204] Mink A.F. *et al* (The ASDEX Upgrade Team) 2018 Nonlinear coupling induced toroidal structure of edge localized modes *Nucl. Fusion* **58** 026011

[205] Trier E. *et al* 2019 ELM-induced cold pulse propagation in ASDEX Upgrade *Plasma Phys. Control. Fusion* **61** 045003

[206] Vanovac B. *et al* 2018 Effects of density gradients and fluctuations at the plasma edge on ECEI measurements at ASDEX upgrade *Plasma Phys. Control. Fusion* **60** 045002

[207] Hansen S.K., Nielsen S.K., Stober J., Rasmussen J., Stejner M., Hoelzl M. and Jensen T. 2020 Parametric decay instabilities near the second-harmonic upper hybrid resonance in fusion plasmas *Nucl. Fusion* **60** 106008

[208] Mink A.F., Wolfrum E., Dunne M., Hoelzl M., Maraschek M., Fischer R., Cavedon M., Harrer G.F. and Stroth U. 2018 Scaling of the toroidal structure and nonlinear dynamics of ELMs on ASDEX Upgrade *Plasma Phys. Control. Fusion* **60** 125011

[209] Orain F. *et al* 2015 Resistive reduced MHD modeling of multi-edge-localized-mode cycles in tokamak *x*-point plasmas *Phys. Rev. Lett.* **114** 035001

[210] Eich T. *et al* (The ASDEX Upgrade Team) 2005 Type-I ELM substructure on the divertor target plates in ASDEX Upgrade *Plasma Phys. Control. Fusion* **47** 815

[211] Park H.K. *et al* 2019 Overview of KSTAR research progress and future plans toward ITER and K-DEMO *Nucl. Fusion* **59** 112020

[212] Kim S.K. *et al* 2020 Nonlinear modeling of the effect of $n = 2$ resonant magnetic field perturbation on peeling-ballooning modes in KSTAR *Nucl. Fusion* **60** 026009

[213] Litaudon X. *et al* 2017 Overview of the JET results in support to ITER *Nucl. Fusion* **57** 102001

[214] Pamela S.J.P., Huysmans G.T.A., Beurskens M.N.A., Devaux S., Eich T. and Benkadda S. 2011 Nonlinear MHD simulations of edge-localized-modes in JET *Plasma Phys. Control. Fusion* **53** 054014

[215] Pamela S. *et al* 2015 Non-linear MHD simulations of ELMs in JET and quantitative comparisons to experiments *Plasma Phys. Control. Fusion* **58** 014026

[216] Zhang W., Cathey A., Hoelzl M., Tierens W. and Bobkov V. 2020 Scattering of ion cyclotron range of frequency waves by filaments and ELMs *Nucl. Fusion* **60** 096001

[217] Labit B. *et al* 2019 Dependence on plasma shape and plasma fueling for small edge-localized mode regimes in TCV and ASDEX upgrade *Nucl. Fusion* **59** 086020

[218] Cathey A. *et al* (The JOREK Team, The ASDEX Upgrade Team, The EUROfusion MST1 Team) 2020 Jorek simulations of multiple type-I ELMs and small ELMs in ASDEX Upgrade *4th Asia Pacific Conf. on Plasma Physics* 26-31, October, 2020 virtual (invited talk)

[219] Harrer G.F. *et al* (The ASDEX Upgrade Team, The EUROfusion MST1 Team) 2020 Stabilizing effects in the AUG small ELM regime *Nucl. Fusion* (private communication)

[220] Garcia-Munoz M. *et al* 2013 Fast-ion redistribution and loss due to edge perturbations in the ASDEX Upgrade, DIII-D and KSTAR tokamaks *Nucl. Fusion* **53** 123008

[221] Garcia-Munoz M. *et al* 2013 Fast-ion losses induced by ELMs and externally applied magnetic perturbations in the ASDEX Upgrade tokamak *Plasma Phys. Control. Fusion* **55** 124014

[222] van Vuuren A.J. *et al* (ASDEX Upgrade Team) 2021 Experimental study of ELM induced fast-ion transport using passive FIDA spectroscopy at the ASDEX Upgrade tokamak *Nucl. Fusion* **61** 046001

[223] van Vugt D.C., Huijsmans G.T.A., Hoelzl M. and Loarte A. 2019 Kinetic modeling of ELM-induced tungsten transport in a tokamak plasma *Phys. Plasmas* **26** 042508

[224] Kirk A. *et al* 2017 Overview of recent physics results from MAST *Nucl. Fusion* **57** 102007

[225] Pamela S.J.P. *et al* 2013 Resistive MHD simulation of edge-localized-modes for double-null discharges in the MAST device *Plasma Phys. Control. Fusion* **55** 095001

[226] Smith S.F., Pamela S.J.P., Thornton A.J., Moulton D., Hoelzl M., Kirk A., Wilson H.R. and Huijsmans G.T.A. (EUROfusion MST1 Team) 2018 Numerical simulations of edge localised modes in MAST-U plasmas *45th European Physical Society Conf. on Plasma Physics (EPS)* (Prague, Czech Republic) p P4.1061 (http://ocs.ciemat.es/EPS2018PAP/pdf/P4.1061.pdf)

[227] Huysmans G.T.A., Pamela S.J.P., Beurskens M.N.A., Becoulet M. and van der Plas E. 2010 Non-linear MHD simulations of natural and pellet triggered ELMs *Proc. of the 23rd IAEA Fusion Energy Conf.* (Daejon, Republic of Korea) p THS/7−1 (https://pub.iaea.org/mtcd/meetings/PDFplus/2010/cn180/cn180_papers/ths_7-1.pdf)

[228] Futatani S., Pamela S., Garzotti L., Huijsmans G.T.A., Hoelzl M., Frigione D. and Lennholm M. 2019 Non-linear magnetohydrodynamic simulations of pellet triggered edge-localized modes in JET *Nucl. Fusion* **60** 026003

[229] Futatani S., Pamela S., Garzotti L., Huijsmans G.T.A. and Hoelzl M. (JOREK Team, JET Contributors) 2019 Non-linear MHD simulations of pellet triggered ELMs in JET *46th European Physical Society Conf. on Plasma Physics (EPS)* (Milan, Italy) pP1.1035 (http://ocs.ciemat.es/EPS2019PAP/pdf/P1.1035.pdf)

[230] Futatani S., Cathey A., Hoelzl M., Lang P.T., Huijsmans G.T.A. and Dunne M. (The JOREK Team, The ASDEX Upgrade Team, and The EUROfusion MST1 Team) 2021 Transition from no-ELM response to pellet ELM triggering during pedestal build-up-insights from extended MHD simulations *Nucl. Fusion* **61** 046043

[231] Lang P.T. *et al* 2014 ELM pacing and high-density operation using pellet injection in the ASDEX upgrade all-metal-wall tokamak *Nucl. Fusion* **54** 083009

[232] Cathey A., Hoelzl M., Futatani S., Lang P.T. and Huijsmans G.T.A. (JOREK Team, ASDEX Upgrade Team, EUROfusion MST1 Team) 2021 Comparing natural and pellet triggered ELMs via non-linear extended MHD simulations *Plasma Phys. Control. Fusion* (in press) https://doi.org/10.1088/1361-6587/abf80b

[233] Artola F.J., Huijsmans G.T.A., Hoelzl M., Beyer P., Loarte A. and Gribov Y. 2018 An in depth look into the physics of ELM triggering via vertical kicks through non-linear MHD simulations *45th European Physical Society Conf. on Plasma Physics (EPS)* (Prague, Czech Republic) http://ocs.ciemat.es/EPS2018ABS/pdf/I2.109.pdf (invited talk I2.109)

[234] Nardon E. 2010 Modélisation non-linéaire du transport en présence d'instabilité MHD du plasma périphérique de tokamak *PhD Thesis* Ecole Polytechnique X (https://pastel.archives-ouvertes.fr/pastel-00003137)

[235] Bécoulet M. *et al* 2014 Mechanism of edge localized mode mitigation by resonant magnetic perturbations *Phys. Rev. Lett.* **113** 115001

[236] Orain F. *et al* 2015 Non-linear MHD modeling of edge localized mode cycles and mitigation by resonant magnetic perturbations *Plasma Phys. Control. Fusion* **57** 014020

[237] Cahyna P. *et al* 2016 Modelling of spatial structure of divertor footprints caused by edge-localized modes mitigated by magnetic perturbations (arXiv:1601.03002)

[238] Thornton A.J., Kirk A., Cahyna P., Chapman I.T., Harrison J.R. and Liu Y. 2014 The effect of resonant magnetic perturbations on the divertor heat and particle fluxes in MAST *Nucl. Fusion* **54** 064011

[239] Orain F. *et al* (The ASDEX Upgrade Team, The EUROfusion MST1 Team) 2017 Non-linear modeling of the plasma response to RMPs in ASDEX Upgrade *Nucl. Fusion* **57** 022013

[240] Orain F. *et al* 2019 Non-linear modeling of the threshold between ELM mitigation and ELM suppression by resonant magnetic perturbations in ASDEX Upgrade *Phys. Plasmas* **26** 042503

[241] Orain F. *et al* (ASDEX Upgrade Team, EUROfusion MST1 Team) 2018 Modeling of the threshold between ELM mitigation and suppression by magnetic perturbations in ASDEX Upgrade *2nd Asia-Pacific Conf. on Plasma Physics (AAPPS-DPP 2018)* (12-17, November 2018 Kanazawa, Japan) (invited oral)

[242] Kim S. 2020 On the mechanism of RMP-driven pedestal transport and ELM suppression in KSTAR *PhD Thesis* Seoul National University (http://dcollection.snu.ac.kr/public_resource/pdf/000000162053_20200929155531.pdf)

[243] Liu F. *et al* 2017 Nonlinear MHD simulations of QH-mode DIII-D plasmas and implications for ITER highQscenarios *Plasma Phys. Control. Fusion* **60** 014039

[244] Liu F., Huijsmans G.T.A., Loarte A., Garofalo A.M., Solomon W.M., Snyder P.B., Hoelzl M. and Zeng L. 2015 Nonlinear MHD simulations of quiescent H-mode plasmas in DIII-D *Nucl. Fusion* **55** 113002

[245] Liu F., Huijsmans G.T.A., Hoelzl M., Viezzer E., Dunne M., Loarte A. and Pamela S. (The ASDEX Upgrade Team, The MST1 Team) 2018 Non-linear MHD simulations of quiescent H-mode in ASDEX upgrade and ITER *Proc. of the 27th IAEA Fusion Energy Conf.* (Gandhinagar, India) p TH/P5−18 (https://nucleus.iaea.org/sites/fusionportal/Shared%20Documents/FEC%202018/fec2018-preprints/preprint0349.pdf)

[246] McAdams R. 2014 Non-linear magnetohydrodynamic instabilities in advanced tokamak plasmas *PhD Thesis* University of York (http://etheses.whiterose.ac.uk/7723/)

[247] Lehnen M. *et al* 2015 Disruptions in ITER and strategies for their control and mitigation *J. Nucl. Mater.* **463** 39−48

[248] Wieschollek F. 2019 Non-linear MHD simulation of mode locking and disruption onset *Master's Thesis* Technische Universität München (http://2.ipp.mpg.de/&tnqx7e;mhoelzl/fwieschollek-master-thesis.pdf)

[249] Sweeney R. *et al* 2018 Relationship between locked modes and thermal quenches in DIII-d *Nucl. Fusion* **58** 056022

[250] de Vries P.C. *et al* 2015 Scaling of the MHD perturbation amplitude required to trigger a disruption and predictions for ITER *Nucl. Fusion* **56** 026007

[251] Nave M.F.F. and Wesson J.A. 1990 Mode locking in tokamaks *Nucl. Fusion* **30** 2575–83

[252] van Overveld T. 2019 Nonlinear JOREK simulations of the Greenwald density limit with a thermo-resistive radiation model *Master's Thesis* Department of Applied Physics, Eindhoven University of Technology

[253] Teng Q., Brennan D.P., Delgado-Aparicio L., Gates D.A., Swerdlow J. and White R.B. 2016 A predictive model for the tokamak density limit *Nucl. Fusion* **56** 106001

[254] Meshcheriakov D. *et al* 2019 Numerical study of tearing mode seeding in tokamak X-point plasma *Phys. Plasmas* **26** 042504

[255] Westerhof E. and Pratt J. 2014 Closure of the single fluid magnetohydrodynamic equations in presence of electron cyclotron current drive *Phys. Plasmas* **21** 102516

[256] Strauss H.R. and Park W. 2000 Pellet driven disruptions in tokamaks *Phys. Plasmas* **7** 250–7

[257] Izzo V.A., Whyte D.G., Granetz R.S., Parks P.B., Hollmann E.M., Lao L.L. and Wesley J.C. 2008 Magnetohydrodynamic simulations of massive gas injection into Alcator C-Mod and DIII-D plasmas *Phys. Plasmas* **15** 056109

[258] Izzo V.A. 2013 Impurity mixing and radiation asymmetry in massive gas injection simulations of DIII-D *Phys. Plasmas* **20** 056107

[259] Izzo V.A. *et al* 2015 The role of MHD in 3D aspects of massive gas injection *Nucl. Fusion* **55** 073032

[260] Ferraro N.M., Lyons B.C., Kim C.C., Liu Y.Q. and Jardin S.C. 2018 3D two-temperature magnetohydrodynamic modeling of fast thermal quenches due to injected impurities in tokamaks *Nucl. Fusion* **59** 016001

[261] Kim C.C., Liu Y., Parks P.B., Lao L.L., Lehnen M. and Loarte A. 2019 Shattered pellet injection simulations with NIMROD *Phys. Plasmas* **26** 042510

[262] Fil A. *et al* 2015 Three-dimensional non-linear magnetohydrodynamic modeling of massive gas injection triggered disruptions in JET *Phys. Plasmas* **22** 062509

[263] Hoelzl M., Hu D., Nardon E. and Huijsmans G.T.A. 2020 First predictive simulations for deuterium shattered pellet injection in ASDEX Upgrade *Phys. Plasmas* **27** 022510

[264] Nardon E., Hu D., Hoelzl M. and Bonfiglio D. 2020 Fast plasma dilution in ITER with pure deuterium shattered pellet injection *Nucl. Fusion* **60** 126040

[265] Wieschollek F., Hoelzl M., Nardon E. *et al* 2021 Shattered pellet injection into an MHD active plasma (private communication)

[266] Izzo V.A. and Parks P.B. 2010 Comment on 'Plasma current spikes due to internal reconnection during tokamak disruptions' *Nucl. Fusion* **50** 058001

[267] Nardon E., Hu D., Hoelzl M., Bonfiglio D., Artola J. and Huijsmans G. (JOREK Team, JET Contributors) 2021 Thermal quench and current profile relaxation dynamics in massive-material-injection-triggered tokamak disruptions *Plasma Physics and Controlled Fusion* (private communication)

[268] Boozer A.H. 2019 Magnetic surface loss and electron runaway *Plasma Phys. Control. Fusion* **61** 024002

[269] Nardon E., Hu D. and Artola J. (The JOREK Team, JET Contributors) 2019 3D non-linear MHD modelling of massive gas injection-triggered disruptions in JET *7th Theory and Simulation of Disruptions Workshop* (Princeton) (https://tsdw.pppl.gov/Talks/2019/Nardon.pdf)

[270] Martín-Solís J.R., Loarte A. and Lehnen M. 2017 Formation and termination of runaway beams in ITER disruptions *Nucl. Fusion* **57** 066025

[271] White R.B. 2001 *The Theory of Toroidally Confined Plasmas* (Singapore: World Scientific)

[272] Strauss H., Joffrin E., Riccardo V., Breslau J., Paccagnella R. and Fu G.Y. 2020 Reduction of asymmetric wall force in JET and ITER disruptions including runaway electrons *Phys. Plasmas* **27** 022508

[273] Bunkers K.J. and Sovinec C.R. 2020 The influence of boundary and edge-plasma modeling in computations of axisymmetric vertical displacement *Phys. Plasmas* **27** 112505

[274] Clauser C.F. and Jardin S.C. 2021 ITER cold VDEs in the limit of perfectly conducting walls *Phys. Plasmas* **28** 012511

[275] Marx A. and Lütjens H. 2017 Hybrid parallelization of the XTOR-2F code for the simulation of two-fluid MHD instabilities in tokamaks *Comput. Phys. Commun.* **212** 90–9

[276] Alenikova K., Huijsmans G.T.A. and Hoelzl M. 2014 Simulations of simplified 3D vertical displacement events in ITER (private communication)

[277] Artola F.J., Huijsmans G.T.A., Lehnen M., Krebs I. and Loarte A. 2019 Halo current studies with self-consistent MHD simulations for ITER 15 MA plasmas *7th PPPL Workshop on Theory and Simulation of Disruptions* (Princeton (USA) 5–7 August 2019)

[278] Gruber O., Lackner K., Pautasso G., Seidel U. and Streibl B. 1993 Vertical displacement events and halo currents *Plasma Phys. Control. Fusion* **35** B191–204

[279] Humphreys D.A. and Kellman A.G. 1999 Analytic modeling of axisymmetric disruption halo currents *Phys. Plasmas* **6** 2742–56

[280] Boozer A.H. 2013 Tokamak halo currents *Phys. Plasmas* **20** 082510

[281] Kiramov D.I. and Breizman B.N. 2017 Model of vertical plasma motion during the current quench *Phys. Plasmas* **24** 100702

[282] Artola F.J., Lackner K., Huijsmans G.T.A., Hoelzl M., Nardon E. and Loarte A. 2020 Understanding the reduction of the edge safety factor during hot VDEs and fast edge cooling events *Phys. Plasmas* **27** 032501

[283] Breizman B.N., Aleynikov P., Hollmann E.M. and Lehnen M. 2019 Physics of runaway electrons in tokamaks *Nucl. Fusion* **59** 083001

[284] Sommariva C., Nardon E., Beyer P., Hoelzl M. and Huijsmans G.T.A. 2018 Electron acceleration in a JET disruption simulation *Nucl. Fusion* **58** 106022

[285] Cai H. and Fu G. 2015 Influence of resistive internal kink on runaway current profile *Nucl. Fusion* **55** 022001

[286] Zhao C., Liu C., Jardin S.C. and Ferraro N.M. 2020 Simulation of MHD instabilities with fluid runaway electron model in M3D-C 1 *Nucl. Fusion* **60** 126017

[287] Sainterme A.P., Sovinec C.R. and Ge W. 2020 Development of a reduced fluid model for runaway electrons in NIMROD simulations *62nd Annual Meeting of the APS Division of Plasma Physics* 9-13 November 2020 virtual p PP12.00035 (https://meetings.aps.org/Meeting/DPP20/Session/PP12.35)

[288] Bandaru V., Artola F.J., Hoelzl M., Papp G., Aleynikov P. and Huijsmans G.T.A. 2018 Non-linear interaction of runaway electrons with resistive MHD modes in an ITER VDE *Proc. of the 27th IAEA Fusion Energy Conf.* (Gandhinagar, India) p TH/Pu–8 (https://nucleus.iaea.org/sites/fusionportal/Shared%20Documents/FEC%202018/fec2018-preprints/preprint0151.pdf)

[289] Bandaru V., Hoelzl M., Reux C., Ficker O., Silburn S., Lehnen M. and Eidietis N. (JOREK Team) 2021 Magnetohydrodynamic simulations of runaway electron beam termination in JET *Plasma Phys. Control. Fusion* **63** 035024

[290] Reux C. *et al* 2021 Runaway electron beam suppression using impurity flushing and large magnetohydrodynamic instabilities *Phys. Rev. Lett.* (private communication)

[291] Dvornova A., Huijsmans G.T.A., Sharapov S., Hoelzl M., Artola J. and Pamela S. 2018 Modelling of TAE mode excitation with an antenna in X-point geometry *45th European Physical Society Conf. on Plasma Physics (EPS)* (Prague, Czech Republic) p P2.1052 (http://ocs.ciemat.es/EPS2018PAP/pdf/P2.1052.pdf)

[292] Dvornova A., Huijsmans G.T.A., Sharapov S., Artola Such F.J., Puglia P., Hoelzl M., Pamela S., Fasoli A. and Testa D. 2020 Modeling of TAE mode excitation with an antenna in realistic X-point geometry *Phys. Plasmas* **27** 012507

[293] Könies A. *et al* (TG ITPA EP) 2018 Benchmark of gyrokinetic, kinetic MHD and gyrofluid codes for the linear calculation of fast particle driven tae dynamics *Nucl. Fusion* **58** 126027

[294] Dvornova A. 2021 Hybrid fluid-kinetic MHD simulations of the excitation of toroidal Alfven eigenmodes by fast particles and external antenna *PhD Thesis* Eindhoven University of Technology and Aix-Marseille University

[295] Strumberger E. and Günter S. 2019 Linear stability studies for a quasi-axisymmetric stellarator configuration including effects of parallel viscosity, plasma flow, and resistive walls *Nucl. Fusion* **59** 106008

[296] Ramasamy R., Hoelzl M., Strumberger E., Nikulsin N., Lackner K. and Guenter S. 2019 Towards linear and non-linear assessments of quasi-axisymmetric stellarators *22nd Int. Stellarator and Heliotron Workshop* (23-27 September 2019 Madison, USA)

[297] Ramasamy R., Hoelzl M., Strumberger E., Yu Q., Lackner K. and (The JOREK Team) 2020 Non-linear MHD stability assessments of stellarators using an axisymmetric approximation *Theory of Fusion Plasmas, Joint Varenna–Lausanne Int. Workshop* (12-16 October 2020 Lausanne, Switzerland)

[298] Ramasamy R., Hoelzl M., Strumberger E., Lackner K. and Günter S. 2021 Nonlinear MHD simulations of external kinks in quasi-axisymmetric stellarators using an axisymmetric external rotational transform approximation *Nucl. Fusion* (private communication)

[299] Beam R.M. and Warming R.F. 1980 Alternating direction implicit methods for parabolic equations with a mixed derivative *SIAM J. Sci. Stat. Comput.* **1** 131–59

[300] Hirsch C. 1991 *Numerical Computation of Internal and External Flows, Volume 1: Fundamentals of Numerical Discretization* (New York: Wiley)