

PAPER • OPEN ACCESS

BLUEMIRA: a modular, open-source framework for designing tokamak fusion reactors

To cite this article: M. Coleman *et al* 2026 *Nucl. Fusion* **66** 126025

View the [article online](#) for updates and enhancements.

You may also like

- [Investigation of high \$Q_{\text{fus}}\$ L-mode plasma operation sustained by elevated pellet fuelling in ITER](#)
J. Zhang, F. Koechl, A.R. Polevoi et al.
- [Experimental investigation of ICRF antenna operation and plasma-wall interaction in the regime of propagating slow waves in front of the antenna](#)
R. Diab, L. Colas, S.G. Baek et al.
- [Nonlinear gyrokinetic simulations of microtearing modes in NSTX and NSTX-U-like plasmas](#)
Tajinder Singh, Tariq Rafiq, Eugenio Schuster et al.

HIDEN
ANALYTICAL

Trusted in Research
for over 40 years



Scan to read the
Physics World Feature

Mass Spectrometry Delivers Fusion Insights

Hidden Analytical featured in **physicsworld**

Discover how Hiden Analytical's specialist mass spectrometry technology, developed through collaboration with Chris Marcus at Oak Ridge National Laboratory, is supporting nuclear fusion research.

H/He Isotope Analysis • Light Gas Separation • Residual Gas Analysis • Fusion Diagnostics

[Click to read the Physics World Feature](#)

www.HidenAnalytical.com

BLUEMIRA: a modular, open-source framework for designing tokamak fusion reactors

M. Coleman^{1,2,3,*} , J.E. Cook¹, F. Franza⁴ , I.A. Maione⁴ , S. McIntosh^{1,5}, H. Baker¹ , A.I. Blair¹, S. Desai¹, O. Funk¹, G. Graham¹, J. Matthews¹, J. Morris¹, C. Mould¹, A. Nilima¹ , H.R.D. Saunders¹, D. Short¹, D. Vaccaro¹  and O. Wong¹

¹ United Kingdom Atomic Energy Authority, Culham Science Centre, Abingdon OX14 3DB, United Kingdom of Great Britain and Northern Ireland

² EUROfusion, Boltzmannstr. 2, Garching 85748, Germany

³ Delta X Fusion S.L.U., Calle del Horno 2, Paracuellos de la Ribera 50299, Spain

⁴ Karlsruher Institut für Technologie, Hermann-von-Helmholtz-Platz 1, Eggenstein-Leopoldshafen 76344, Germany

⁵ Science Division, ITER Organisation, Route de Vinon-sur-Verdon, St. Paul-lez-Durance, CS 90 046; 13067, France

E-mail: matti.coleman@ukaea.uk

Received 17 June 2026

Accepted for publication 10 September 2026

Published 28 September 2026



Abstract

As efforts across the world to deliver fusion power enter their respective conceptual design stages, many begin to encounter the difficulties of generating complete designs, evaluating them, performing trade-off studies, and comparing different alternatives. Many organisations successfully use fusion reactor systems codes with low-fidelity formulations of the critical physics and engineering aspects to inform the initial design stage and steer them towards a region of the design space they consider attractive, but struggle to explore the design space in more depth. This work is an overview of the open-source tokamak reactor design framework, BLUEMIRA, built from the BLUEPRINT (Coleman and McIntosh 2019 *Fusion Eng. Des.* **139** 26–38; Coleman 2021 *Thesis*, University of Cambridge) and MIRA (Franza 2019 Development and validation of a computational tool for fusion reactors system analysis; Franza F. *et al* 2022 *Nucl. Fusion* **62** 076042) codes, which goes beyond the ‘classical’ systems code paradigm and enables the user to parameterise their reactor designs at higher fidelity. Reactor design workflows ranging from 0-D ‘radial build’ design, equilibrium design, to 3-D CAD and automated neutronics analyses can be run within minutes. We present the functionality implemented in the BLUEMIRA code and demonstrate its application to the conceptual design of a conventional aspect ratio fusion reactor, performing an indicative design study over aspect ratio and number of toroidal field coils.

* Author to whom any correspondence should be addressed.



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

Keywords: tokamak, systems code, design optimisation, parameterisation

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

1. Introduction

As efforts across the world to deliver fusion power enter their respective conceptual design stages, many begin to encounter the difficulties of generating complete designs, evaluating them, performing trade-off studies, and comparing different alternatives. Many organisations, commercial or otherwise, successfully use fusion reactor systems codes with low-fidelity formulations of the critical aspects to inform the initial design stage and steer them towards a region of the design space they consider attractive.

Once there, however, progress has a tendency to slow down. The nature of the problem inherently requires a broad set of multi-disciplinary capabilities in order to be addressed coherently. Different people using different tools must work in unison on a single idea. Assumptions, geometry, load cases, boundary conditions, and the like must be exchanged and evaluated separately and yet consistently. Many studies depend on the results of other studies. Results must be post-processed, analysed, and coherently combined, ultimately leading to informed changes to the initial design.

Depending on the organisation, its staff and capabilities, and the funding it controls, the process of changing a fusion reactor design can be slow. Depending on the change, dozens of domain experts may be called upon to perform analysis or design activities using different geometry, plasma scenarios, load cases, etc. Few are the tokamak systems that can be designed in isolation. Typically, many weeks or months are necessary before a design change can be properly evaluated and a decision made on whether or not to include it in the design point. Significant or more fundamental changes, e.g. alternative magnet concepts or changes to key machine parameters, such as aspect ratio, can languish in study for years.

More often than not, a fusion reactor design effort will centre its focus on a single reactor design and, at best, perform a few high-level trade-off studies around that point using 0/1-D systems codes. As a result, the conceptual design phase typically involves a small number of incremental iterations around an initial design point studied in depth, rather than a systematic exploration of the broader parameter space and an equally thorough investigation of alternative concepts. In a field as uncertain as nuclear fusion, and with designs as complicated as tokamaks, this is hardly an ideal situation.

Several authors have made the same observations and proposed frameworks to address it: MIRA [1, 2], BLUEPRINT [3, 4], FUSE [5, 6], and FREDa [7]. These codes were designed to accelerate the fusion reactor design process, combining higher fidelity models and parameterised routines to develop and analyse fusion reactor concepts faster and in more depth. More organisations may be working on similar frameworks, but as many of the best-funded fusion programmes are now

commercial in nature, they may not be willing or at liberty to publish their work.

In this work, we present our take on addressing this problem: the tokamak reactor design framework BLUEMIRA, which is the result of a collaboration between UKAEA and KIT, combining BLUEPRINT and MIRA. Both these codes had a similar starting point (a radial build from a traditional systems code) and a similar scope (designing a reactor in more detail in 2-D and 3-D, performing a variety of typical tokamak design activities). They also were both primarily developed for EU-DEMO design studies, and thus were well-suited for amalgamation.

We make the following distinction here between ‘systems code’ and ‘reactor design framework’: a systems code solves a pre-defined reactor optimisation problem self-consistently, returning an optimum. A reactor design framework may be used to construct parameterisations of reactor concepts, and then used to incrementally modify certain aspects in response to evolving demands. The former is more commonly used for design exploration and sensitivity studies, the latter is more commonly used for to rapidly perform design iterations on an existing concept. Whilst a systems code may be to some extent used as a ‘black-box’, a reactor design framework includes engineers in the optimisation process; giving them the building blocks to parameterise their design, and proceduralise their design process.

We first discuss the general use cases and requirements for a tokamak reactor design framework. In the second section, we introduce the BLUEMIRA code, describing it in broad terms. In the third section we describe some of the individual modules in more detail, and in the fourth section we describe the nature of design procedures in BLUEMIRA by means of an example. Finally, in sections five, six, and seven we demonstrate the application of the BLUEMIRA code in a series of case studies on a particular reactor design problem: a conventional aspect ratio power reactor, before discussing our work and presenting our conclusions.

1.1. Fusion reactor conceptual design phases

At present, globally, a wide range of conceptual fusion reactor studies are underway. Many publicly-funded research bodies are working on fusion reactor concepts, such as EUROfusion with EU-DEMO [8], UKAEA/UKIFS with STEP [9], various Chinese institutions with CFETR [10], various Japanese institutions with JA DEMO [11], NFRL with K-DEMO [12], and a variety of studies in the US, e.g. [13, 14]. These studies are continuing a body of historical conceptual fusion reactor design studies in the field, most of which focus on tokamaks, see [15]. Over the last decade, a number of notable privately-funded entities have also embarked on their own fusion reactor design studies, such as Commonwealth Fusion Systems with ARC [16, 17] and Tokamak Energy [18]. Some of these private

entities are spin-offs from national laboratories, and continue to work alongside publicly-funded institutions by means of collaboration agreements or public-private partnerships. The majority of these fusion start-ups focus on magnetically confined fusion, with the more established focussing on tokamaks or stellarators.

These studies are, without exception, so-called ‘academic reactors’ [19], and should be treated as such.

The conceptual design phases of fusion power reactors presently underway in the world have a lot in common: (i) they are plagued with high uncertainties, in particular in terms of plasma performance; (ii) they rely upon a cocktail of non-existent and/or low-maturity technologies; (iii) relatively few design decisions, technological choices, and parameters are set in stone.

1.2. Requirements for a tokamak fusion reactor design framework

Here we outline our views on what is required for a tokamak fusion reactor design framework.

The overall mission of a tokamak fusion reactor design framework, as we see it, is to support physicists and engineers carrying out the conceptual and engineering design of tokamak fusion reactors.

Based on observations of fusion reactor design efforts, we foresee the following use cases:

- (i) Generation of conceptual tokamak reactor design points
- (ii) Modification of an existing reactor design point
- (iii) Exploration of the design space of a particular reactor concept
- (iv) Comparison of alternative sub-system and reactor concepts

For any design framework to be as useful as possible, to as many people as possible, the following aspects are important:

- (i) Ease of access (licence, installation, platform dependency, etc)
- (ii) Ease of use (application programming interface (API) design, language, etc)
- (iii) Ease of understanding (documentation, tutorials, examples, training, etc)
- (iv) Ease of interoperability (file I/O, unit conventions, visualisation, etc)
- (v) Robustness
- (vi) Verification and validation

For tokamak design in particular, we consider the following features as necessary components of a framework:

- (i) 0/1-D systems code solver
- (ii) 1.5-D plasma transport and equilibrium solver
- (iii) 2-D magneto-hydrodynamic (MHD) equilibrium solver(s)
- (iv) 3-D magnetostatics solver(s)
- (v) 2-D and 3-D geometry creation and manipulation

- (vi) 2-D and 3-D radiation transport solver(s)
- (vii) 2-D and 3-D finite element solver(s)
- (viii) Material properties
- (ix) Fuel cycle model
- (x) Power balance model
- (xi) Cost model

In terms of fidelity and run-time performance, it is vital that a balance is struck. Users must be able to perform a run quite quickly, that is to say, within a matter of minutes. However, they must also be able to perform runs at variable levels of fidelity, such that when one hones in on a particular design point, it can be run at higher fidelity, with commensurate increases in run-time.

Furthermore, the functionality offered should be generic, such that it is applicable to as many tokamak reactor designs as possible, and modular, such that somebody interested specifically in e.g. plasma equilibria, need not interact with the fuel cycle or cost models.

2. The BLUEMIRA framework

2.1. Status

In 2020, active development of the BLUEPRINT and MIRA codes ceased, and since then some authors of these codes, and many more people, have been working on the development of a new fusion reactor design framework: BLUEMIRA.

BLUEMIRA version 1 was released in early 2023, and is presently at version 2.15.0. A significant effort has gone in to creating a stable software architecture that is capable of supporting long-term modular development by multiple people in different organisations. The architecture has been designed to be less specific to any one tokamak concept, and easier to use for non-experts.

BLUEMIRA is actively used on the EU-DEMO reactor design [20], and in the STEP reactor design⁶ [21–24].

2.2. Software aspects

2.2.1. Licence. If commercial fusion power is ever to be achieved, it will likely be the result of many decades of research, most of which will not have been performed in the entities which are ultimately successful. In the hope that it may be useful for the community as a whole, and in the interests of improving transparency in fusion research, we have chosen to release BLUEMIRA open-source [25]⁷ distributed under the LGPL 2.1+ licence [26].

LGPL 2.1+ is a weak copy-left, open-source licence that allows for: commercial use, modification, distribution, and private use. This licence does not provide a warranty or offer liability. The code is publicly available and leverages a variety of open-source tools to perform automatic code functional and

⁶ The STEP team contributes regularly to the open-source BLUEMIRA repository (generic features, issues, etc), but implements their reactor design parameterisation in a separate, closed-source repository.

⁷ <https://github.com/Fusion-Power-Plant-Framework/bluemira>.

quality checks, to host the BLUEMIRA documentation⁸, and to provide issue and code management.

Making the code open-source in this way allows external collaborators to access and potentially contribute to the code while ensuring that the core development team has oversight on requirements, bugs, and priorities. The team makes use of code review on each change to allow comment on how a change or fix has been implemented, and to uphold the programming and numerical standards, as well as conventions.

2.2.2. Language. The choice of programming language for any program is an important one as it can greatly affect the ease of use, the speed of development, and the type and quantity of collaborators. Python 3 was chosen for a few reasons. As an open-source, dynamically-typed language with a very simple syntax, the barrier to entry for engineers and physicists not necessarily trained in computer science is quite low. The ease of use and speed of development are relatively high, making it an ideal language for prototyping functionality and architecture, as well as attracting new users and developers. Importantly, a wide variety of open-source modules and packages are available as dependencies, enabling a broad set of functionality to be leveraged incurring minimal maintenance overheads.

Furthermore, BLUEPRINT was written in Python 3 therefore only some parts of MIRA needed to be translated (from MATLAB®, which is not an open-source language) in order to combine the two codes.

Another benefit of using Python is the ease with which other programs or languages can be called, enabling greater collaboration and integration with other programs.

2.2.3. Platform. Linux was selected as the main supported operating system for three main reasons:

- (i) It is free and open-source. This facilitates distributed computing on clusters in future.
- (ii) Some dependencies (e.g. OpenMC [27]) are not available natively in Windows.
- (iii) Linux is available quasi-natively on Windows in the form of Windows Subsystem for Linux (WSL2), which facilitates access on Windows machines.

BLUEMIRA is only tested on Ubuntu, but we have both users and developers who work with the code on Windows via WSL2, and natively on Mac OS X.

Work is ongoing to set up a Docker [28] image of BLUEMIRA to further facilitate access and ease of use.

2.3. Conventions

2.3.1. Units. The SI unit convention is used throughout the BLUEMIRA framework. Values can be input with associated non-SI units, and are automatically converted to SI. All




Figure 1. The three different coordinate systems used in the BLUEMIRA framework: right-handed Cartesian coordinate system (x, y, z , in black), cylindrical coordinate system (x, ϕ, z , in blue), and quasi-toroidal coordinate system (R, r, θ, ϕ , in red).

outputs are in SI, but can be trivially converted. This is particularly important when interfacing with other codes, for which unit mappings exist to facilitate I/O.

User inputs for angles are specified in degrees but integral trigonometric functions and periodic evaluation functions use radians with an automatic conversion at the API boundary.

2.3.2. Coordinate systems. Due to the variety of different activities carried out in BLUEMIRA, several coordinates systems are required, see figure 1.

In BLUEMIRA, given the need for multiple coordinate systems, x is used to describe the radial coordinate, traditionally denoted r , even in 2-D cylindrical coordinates. This facilitates the transition from 2-D to 3-D coordinate systems.

2.3.3. Structure. The aim of BLUEMIRA is not to reinvent the wheel but provide an open-source framework in which many tools and features useful for tokamak design can be used in isolation or combined together relatively seamlessly.

BLUEMIRA comprises several modules, which are relatively independent of one another. Broadly speaking, there are two categories of module: support modules, and discipline modules, see figure 2.

The design intent behind this modularity is simple: if one is only interested in a particular aspect, say the tritium fuel cycle, or magnets, then one only need interact with the corresponding modules. If, on the other hand, one wants to design a whole reactor, then the tools are available in one place.

The ‘codes’ module contains interfaces to key external codes. Particular effort has been spent on abstracting out the interfaces to these codes, such that it is possible to replace these dependencies at a later date, should it prove necessary or useful to do so, without disrupting the user-facing API.

Another key motivation behind abstract interfaces for external codes, is that it makes it possible to use an alternative tool for a specific activity. For example, if a user would prefer to use a systems code other than PROCESS, they could write an interface to it using the mechanics in the ‘codes’ module.

⁸ <https://bluemira.readthedocs.io/en/latest/>.




Figure 2. Schematic of the various BLUEMIRA modules.

In some cases, we have opted to split out relatively self-contained pieces of functionality into separate repositories: an EQDSK [29] reader and COCOS [30] conversion tool, eqdsk [31], a materials property module and library, matproplib [32], a highly representative tokamak neutron source, tokamak-neutron-source [33], and a DAGMC [34] geometry pre-processing pipeline, fast-ctd [35]. These will be introduced in this work alongside the main BLUEMIRA functionality. The motivation for separating these modules is that it is often easier for external open-source collaborators to engage with a tool that has a small, generic, and well-defined purpose. We hope that these tools will be useful to the wider open-source scientific community, whom we hope will also help to improve and maintain them.

2.3.4. File I/O. The JavaScript Object Notation (JSON) Data Interchange Format [36] is used for all input files. All outputs can also be saved as JSON files, although in some cases more familiar file standards are more commonly used (e.g. EQDSK and IMAS [37] for plasma equilibria, STP [38] for CAD objects).

3. BLUEMIRA modules

In this section, the key BLUEMIRA modules are described individually.

3.1. 0/1-D systems code: PROCESS [39, 40]

BLUEMIRA uses PROCESS [39, 40] as its ‘radial build solver’; in other words, it is the engine that determines a design point at low fidelity (0/1-D).

PROCESS is used to solve an optimisation problem in which an objective function is minimised, subject to various bounds and constraints on the design space.

The variable vector, variable bounds, objective function, and constraint functions and values are all chosen by the user, making it quite a generic tool. From any given PROCESS run, the radial build (see figure 3) and several dozen values for various parameters are extracted and used throughout the rest of




Figure 3. Graphical representation of a typical PROCESS radial build.

the BLUEMIRA run. Some of these are taken as ‘gospel’, others are overwritten by the results of higher fidelity models.

The PROCESS team is presently developing a modernised data structure and Python 3 API which BLUEMIRA will leverage as soon as it is ready.

3.2. 1.5-D plasma transport and fixed boundary equilibrium: PLASMOD [41, 42]

BLUEMIRA has an interface to the PLASMOD code [41, 42], which uses the 2-D fixed boundary equilibrium code EMEQ [43, 44] to compute the equilibria to update terms in the solution of the 1-D current diffusion equation. Thus, it solves the two equations with consistent pressure and current profiles, in a so-called 1.5-D approach, common in the field. EMEQ is a three-moment equilibrium solver which uses an up-down symmetric parameterised description of the flux surfaces (i.e. no X-points).

The PLASMOD code has a variety of models, options, and constraints that can be used. These include, but are not limited to, different transport models, different scaling laws (e.g. pedestal, L-H transition), whether or not to use I_p or q_{95} as an input, whether or not to fixed fusion power or H-factor, divertor protection constraints, etc.

It can be used to obtain certain plasma performance parameters (e.g. P_{fus}) at higher fidelity than PROCESS. PLASMOD also outputs the temperature, density, current, and pressure profiles, as well as the flux functions p' and FF' which are useful for subsequent equilibrium calculations.

As PLASMOD is not open-source, alternative implementations of the 1.5-D fixed boundary equilibrium and transport




Figure 4. Semi-analytic functions for poloidal magnetic flux, radial field, and vertical field due to a rectangular current source of uniform current density in cylindrical coordinate axisymmetry.

paradigm are being investigated, with the aim of incorporating an interface in BLUEMIRA to a publicly available external code that performs such calculations. Ideally, such a module would include the capability of handling an up-down asymmetric plasma boundary, thereby simplifying the process of making a fixed boundary equilibrium consistent with a free boundary equilibrium.

BLUEMIRA can nevertheless be run without PLASMOD, taking plasma profiles and integral quantities from PROCESS, transport solvers such as JETTO [45], or user input.

3.3. Magnetostatics

The magnetostatics module contains some useful functions and objects to calculate the magnetic fluxes and fields from a range of different 2-D and 3-D current sources, analytically and semi-analytically. The fact that the calculations are analytical and semi-analytical mean that the fields can be evaluated at arbitrary points in space without resorting to the specification of a mesh or finite element problem. This is particularly important for run-time performance, and even more so when using the field calculations in objective or constraint functions within optimisation problems.

In cylindrical axisymmetry about the z -axis, Green's functions for the poloidal magnetic flux ψ , its partial derivatives ($\frac{\partial \psi}{\partial x}$ and $\frac{\partial \psi}{\partial z}$), and the radial and vertical magnetic field, B_x and B_z , (widely available in the literature), can be used for current point sources.

Poloidal field coils can be modelled in 2-D using quadratures of Green's functions, but can also be modelled as rectangular current sources of uniform current density, following semi-analytic calculations for B_x and B_z in cylindrical coordinates given in [46], see figure 4. ψ is obtained here by numerical integration.

In 3-D Cartesian coordinates, the magnetostatics module offers a simple filamentary Biot-Savart current source (see equation (1)), and two 3-D uniform current density current sources of rectangular cross-section: trapezoidal prisms (see

[47, 48]) and circular arcs (see [49]), see figure 5, using semi-analytical reductions of equation (2),

$$\mathbf{B}(\mathbf{r}) = \frac{\mu_0 I}{4\pi} \int_{\mathcal{L}} \frac{d\mathbf{l} \times (\mathbf{r} - \mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|^3} \quad (1)$$

$$\mathbf{B}(\mathbf{r}) = \frac{\mu_0}{4\pi} \iiint_{\mathcal{V}} \frac{\mathbf{J}(\mathbf{r}') \times (\mathbf{r} - \mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|^3} d\mathbf{r}' \quad (2)$$

where:

$d\mathbf{l}$ = differential vector of a current-carrying path

$\mathbf{r}$ = sink point $\{x, y, z\}$ at which the field is to be calculated

$\mathbf{r}'$ = source point $\{x', y', z'\}$ containing the centre of the current source

$\mathcal{L}$ = length domain where the current flows

$d\mathbf{r}'$ = differential volume of a current-carrying element

$\mathbf{J}$ = current density vector

$\mathcal{V}$ = volume domain where $\mathbf{J}$ is defined

The reader is referred to [47, 48] for a complete description of the equations used to calculate magnetic fields from 3-D current sources from trapezoidal prisms of uniform current density, and to [49] for circular arcs of rectangular cross-section and uniform current density. These equations are solved by numerical integration (and are more tractable than solving equation (2) directly), and are similar to those used in the closed-source, no-licence, legacy EFFI code [50], which provides formulations for a circular arc, but only treats the simpler case of a cuboid rather than a trapezoid.

Importantly, the trapezoidal current sources can be assembled with different end angles, enabling relatively complex paths to be discretised without overlapping sources. Utilities are provided to assemble current sources from ordered sets of planar coordinates describing the current centreline (CCL) of a rectangular cross-section coil, see figure 6. This is vital for modelling TF coils, and accurately and rapidly calculating the peak field in the TF coil winding pack (including 3-D effects).




Figure 5. Biot-Savart, trapezoidal prism, and circular rectangular cross-section current sources shown in blue, with magnetic field $|B|$ shown as filled contours in 3-D Cartesian coordinates.




Figure 6. Trapezoidal prism and circular arc of rectangular cross-section current sources shown in blue, for a 20 MA circular coil of radius 6 m, breadth 0.5 m, depth 1 m, centred at radius from the z -axis of 9 m, patterned symmetrically about the z -axis, with magnetic field $|B|$ shown as filled contours in 3-D Cartesian coordinates. The trapezoidal representation uses a discretisation of the circle to create 25 current sources, whereas the circular representation is a single current source.

3.4. Equilibria

Two different ideal plasma magnetohydrodynamic (MHD) 2-D axisymmetric equilibrium solvers are included in the BLUEMIRA code: a fixed boundary finite element solver and a free boundary finite difference solver. Both solvers compute the equilibrium configuration by solving a reduced form of Maxwell's equations for the poloidal magnetic flux in axisymmetry about the z -axis that, in its general form, reads:

$$\Delta^* \psi = -\mu_0 x J_\phi \quad (3)$$

where:

$$\Delta^* \psi = \frac{\partial^2 \psi}{\partial z^2} + x \frac{\partial}{\partial x} \left(\frac{1}{x} \frac{\partial \psi}{\partial x} \right)$$

ψ = poloidal magnetic flux

$\mu_0 = 4\pi \times 10^{-7}$ is the permeability of vacuum

J_ϕ = toroidal current density

In the plasma region, under the assumptions of axisymmetry and ideal MHD, this formulation reduces to the well-known Grad-Shafranov equation [51, 52], which expresses

the force balance between magnetic and pressure forces in the form:

$$\Delta^* \psi = -\mu_0 x^2 \frac{dp}{d\psi} - \frac{1}{2} \frac{dF^2}{d\psi} \quad (4)$$

where $p(\psi)$ is the plasma pressure, and $F(\psi) = R_0 B_\phi$ relates to the toroidal field.

Equation (3) forms the common basis for both solvers. In the fixed boundary case, the toroidal current density, J_ϕ , is prescribed within a known plasma domain, and no external sources are included. Conversely, in the free boundary formulation, J_ϕ includes both the plasma and external contributions (e.g. from poloidal field coils), and the plasma boundary is not imposed but instead obtained self-consistently from the solution.

The fixed boundary solver is used primarily for detailed design-point studies, as well as for equilibrium reconstruction or for benchmarking numerical solvers, where the separatrix or last closed flux surface (LCFS) is provided from design, experiment, or MHD codes. Regarding the latter, in BLUEMIRA, the fixed boundary solver can be used to obtain a self-consistent

up-down asymmetric equilibrium solution with profiles from a PLASMOD run with an inherently up-down symmetric boundary representation.

In contrast, the free boundary solver is typically employed in the poloidal coil design workflow (e.g. optimisation of the coils' current and position), where the plasma boundary emerges self-consistently from the interaction between the plasma current distribution and external conductor systems.

3.4.1. Fixed boundary solver. To determine the equilibrium, the Grad–Shafranov equation is solved assuming a known plasma boundary. The solution domain Ω is bounded by a predefined contour, where Dirichlet boundary conditions are imposed: $\psi = \psi_b$. Geometry and mesh generation are performed using gmsh [53], and the finite element discretisation is implemented using the Fenics [54] library and its modern backend DOLFINx [55].

The weak form of the problem is given by: find $\psi \in V$ such that

$$\int_{\Omega} \frac{1}{x} \nabla \psi \cdot \nabla \mathbf{v} d\Omega = \int_{\Omega} \left(-\mu_0 x \frac{dp}{d\psi} - \frac{1}{2x} \frac{dF^2}{d\psi} \right) \mathbf{v} d\Omega \quad \forall \mathbf{v} \in V_0, \quad (5)$$

where V is the finite element space of continuous Lagrange functions and $V_0 \subset V$ vanishes on the boundary.

A critical point is that the profiles p' and FF' are usually not directly expressed as functions of ψ , but instead depend on the normalised flux, usually given as:

$$\psi_n(\psi) = \frac{\psi_{ax} - \psi}{\psi_{ax} - \psi_b}, \quad (6)$$

where ψ_{ax} is the value of ψ at the magnetic axis, defined as the maximum of ψ within Ω . Since ψ_{ax} is only known *a posteriori*, a fully consistent non-linear formulation cannot be directly passed to the standard DOLFINx non-linear solver interface.

To handle this, a Picard-like fixed-point iteration is adopted. At each iteration, the current solution ψ^k is used to compute $\psi_{ax}^k = \max_{\Omega} \psi^k$, and from it the normalised flux $\psi_n^k(\psi^k)$ is constructed. The profiles $p'(\psi_n^k)$ and $FF'(\psi_n^k)$ are then evaluated and projected onto the finite element space. This yields a linear elliptic problem with fixed source terms, which is solved to obtain ψ^{k+1} . Convergence is monitored through the relative change in the solution norm. Under-relaxation may optionally be applied to improve stability, particularly for steep profile gradients.

The variational problem at each iteration reads: find $\psi^{k+1} \in V$ such that

$$\begin{aligned} & \int_{\Omega} \frac{1}{x} \nabla \psi^{k+1} \cdot \nabla \mathbf{v} d\Omega \\ &= \int_{\Omega} \left(-\mu_0 x p'(\psi_n^k) - \frac{1}{2x} FF'(\psi_n^k) \right) \mathbf{v} d\Omega \quad \forall \mathbf{v} \in V_0. \end{aligned} \quad (7)$$

This formulation supports arbitrary smooth profile functions, including experimentally fitted or tabulated ones. Once

convergence is reached, post-processing includes the extraction of magnetic surfaces, final current density profiles, and magnetic field components. These are exported in standard formats for further interface with other BLUEMIRA modules.

Figure 7 shows a verification test of the fixed boundary equilibrium solver with respect to an analytic Solov'ev equilibrium found in [56].

3.4.2. Free boundary solver. The free boundary equilibrium solver is described in [4, 57], appendix A. It is an inverse, stationary solver, based on the finite difference approach and uses Picard iteration (with relaxation) to solve for equilibria under certain user-defined conditions. The free flux functions p' and FF' (as a function of the normalised flux, see equation (6)) can be tailored to match certain plasma integrals (I_p , β_p , and l_i), taken as user-inputs, or from the output of a 1.5-D fixed boundary equilibrium and transport solution.

Equation (3) is solved with external coil currents:

$$J_{\phi} = J_{\phi,pl} + J_{\phi,coils}. \quad (8)$$

A variety of different convergence criteria are available for use, depending on the user's intent.

The principal use-case of the free boundary equilibrium solver is for poloidal field (PF) system design. As such, BLUEMIRA contains a relatively flexible set of equilibrium constraints which can be applied to specific points or regions in an equilibrium and PF system optimisation problem. These include:

- (i) flux: $\psi(x, z)$,
- (ii) isoflux: $\psi(x, z) = \psi(x_{ref}, z_{ref})$,
- (iii) radial field: $B_x(x, z)$,
- (iv) vertical field: $B_z(x, z)$,
- (v) spherical harmonics [58, 59],
- (vi) toroidal harmonics [60, 61],

The equilibrium constraints take the following form:

$$\mathbf{GI} - \mathbf{b}_t + \mathbf{b}_p = \mathbf{0} \quad (9)$$

where:

$\mathbf{G} = n_t \times n_c$ matrix of unit current responses at the control points

$\mathbf{I} = n_c$ vector of coil currents

$\mathbf{b}_t = n_t$ vector of desired target values

$\mathbf{b}_p = n_t$ vector of passive, uncontrolled current contributions (including the plasma)

These are all linear, and can be used either as equality constraints, inequality constraints, or in the objective function, and used to reach, e.g. a desired LCFS shape and position. A variety of optimisation objectives are also supported, including L_2 -norm minimisation of the currents, target errors, and a Tikhonov regularisation of the currents with the L_2 -norm of the target errors, see equation (10).




Figure 7. (left) Fixed boundary equilibrium solution (right) absolute error compared with the Solov'ev analytical solution.

$$\begin{aligned} & \|\mathbf{I}\|_2^2 \\ & \|\mathbf{GI} - \mathbf{b}_t + \mathbf{b}_p\|_2^2 \\ & \|\mathbf{GI} - \mathbf{b}_t + \mathbf{b}_p\|_2^2 + \|\mathbf{\Gamma I}\|_2^2 \end{aligned} \quad (10)$$

where $\mathbf{\Gamma}$ is the regularisation term, ordinarily the identity matrix times some low factor ($\approx 1 \times 10^{-8}$), depending on the user's intent.

Typically the regularised norm of the target errors to some desired plasma shape is used to solve for coils currents to obtain an initial inverse equilibrium. Figure 8 shows a benchmark with the *FreeGSNKE* code [62], solving an equivalent unconstrained inverse problem for a typical MAST-U equilibrium using the same set of coils on the same grid, minimising the error to the same set of constraints with Tikhonov regularisation. The modest discrepancies are largely due to subtle differences in implementation and default values (e.g. for the regularisation term, which ultimately means that the optimal current vectors differ).

Crucially for the design of PF systems, current bounds can be specified, as well as vertical force F_z and poloidal field B_p constraints for the coils are also available to the user, see [57].

Note that these constraints on the coils are not linear, and non-linear optimisation methods must be used in order to solve such problems. Typically, this involves solving sequential quadratic sub-problems in which the constraints are linearised. In *BLUEMIRA*, this is typically done using either the *SLSQP* [63] or *COBYLA* [64] constrained optimisation algorithms.

Further non-linear constraints on the flux expansion, f_x , connection length, L_c , and grazing angle, α_g , are now also available, to support the more detailed design of diverted equilibria.

Two approaches for the optimisation of PF coil positions are also supported: a position 'track' for the PF coil centroids,

as described in [57], as well a set of convex region constraints for the PF coil centroids. PF coil sizes are dictated by user-input current densities. The position optimisation of the PF coils is performed separately to the optimisation of the currents, where the solution vector is the positions of the PF coils. The positions are constrained such that the coils lie on or within the specified regions.

See section 5.4.3 for a more in-depth application of constrained coil current and position optimisation using the free boundary equilibrium solver.

3.4.3. Vertical stability considerations. An implementation of the *RZIp* code [65, 66] based on [67] was recently added to *BLUEMIRA*, and has been successfully benchmarked with the *FIESTA* [68] code's implementation. This is used to assess the passive vertical stability of an equilibrium, given the presence of some toroidally continuous, passively conducting structures surrounding the plasma (e.g. vacuum vessel). *RZIp* gives a simple metric for instantaneous stability of the plasma under a small perturbation, f_s (see equation (11)) [67]. The plasma current vector, $\mathbf{I}_p$, is calculated as the product of the toroidal current density of the plasma over the grid and the cell size used in the equilibrium.

$$f_s = \frac{\mathbf{I}_p^T \mathbf{M}'_{p,s} \mathbf{M}_{s,s}^{-1} \mathbf{M}'_{s,p} \mathbf{I}_p}{\mathbf{I}_p^T \mathbf{M}''_{p,a} \mathbf{I}_a} \equiv -\frac{F'_s}{F'_d} \quad (11)$$

where

$\mathbf{I}_k$ = Current vector of conductor group k

$\mathbf{M}_{i,j}$ = Mutual inductance matrix between conductor group. i and j

$\mathbf{p}$ = plasma conductor group

$\mathbf{s}$ = passive conductor group

$\mathbf{a}$ = active conductor group




Figure 8. MAST-U static inverse free-boundary equilibrium as solved in BLUEMIRA, showing ψ as contour lines, and the separatrices as dashed blue lines. The FreeGSNKE [62] separatrices are shown in red. The blue rectangles represent the coils and the orange circles and crosses represent isoflux and field null ‘targets’, respectively, which were used in both codes.

$$l = \frac{d}{dz}$$

z = vertical coordinate

F'_s = stabilising force gradient

F'_d = destabilising force gradient

It is possible to include the stabilising effect of active systems in f_s by including the inductances of the active systems in the calculation of $\mathbf{M}_{s,s}$ and $\mathbf{M}'_{p,s}$ and its transpose [67].

The stability parameter f_s is often defined in terms of the margin of stability, given by $m_s = f_s - 1$. Due to $f_s < 1$ being undefined [67], m_s is used as a more intuitive metric.

The margin of stability is proportional to the growth rate of the vertical movement in the plasma, γ_s , through the

inductance-resistance time constant, τ_{passive} , of the passive systems [67],

$$\gamma_s = \frac{1}{\tau_{\text{passive}}} \frac{1}{m_s}. \quad (12)$$

Inclusion of simplified plasma stability models in BLUEMIRA enables more in-depth analysis of tokamaks in the early concept design stage.

3.5. Geometry

Parameterised geometry is one of the key features for a reactor design framework. MIRA had parameterised 2-D geometry in discretised form (that is to, say using many points to describe a curve) and BLUEPRINT had parameterised geometry in 3-D, also in discretised form, with some utilities to convert some of the point clouds to splines when exporting the CAD to files. Perhaps the most significant improvement implemented in BLUEMIRA with respect to its predecessors was the move to ‘primitive’ descriptions of geometry (boundary representations) from the outset. The idea is that holding accurate descriptions of geometry in terms of primitive CAD types (rather than clouds of points) eases the transition to meshing, and the specification of finite element problems and radiation transport problems. When using primitives one can label individual surfaces and pass these labels to finite element problems in order to support the parametric specification of boundary conditions and loads.

A decision was made to use FreeCAD [69] as the underlying CAD engine, which itself wraps OpenCASCADE [70], which is the only fully-featured open-source CAD package available at present. A wrapper around the FreeCAD functionality is used to protect users against possible changes of the underlying CAD engine by minimising any eventual API changes should BLUEMIRA choose to use another CAD engine in future.

The typical hierarchy of CAD types is exposed in BLUEMIRA, with the terminology following that of OpenCASCADE, namely: wire $\rightarrow$ face $\rightarrow$ shell $\rightarrow$ solid $\rightarrow$ compound. Similarly, classical CAD creation operations such as extrude, revolve, sweep, loft, Boolean cut and fuse are exposed to the user, along with the various CAD transformation operations, such as translate, rotate, mirror. Geometry objects can also be saved and exported into a wide variety of common CAD file formats.

Using these operations, it is relatively trivial to proceduralise the creation of geometry, and parametrise designs such that CAD models of even complex parts can produced in seconds, see e.g. Shimwell *et al* [71] who proceduralise the CAD generation of relatively detailed breeding blanket module geometries. As the creation of CAD models is one of the key bottlenecks in engineering design, it is vital that this is not done manually when exploring ranges of design variables.

To support the optimisation of geometry, several geometry parameterisations are made available in BLUEMIRA. These geometry parameterisations encapsulate wire shapes typical in tokamak design, such as ‘D’ shapes used for TF coil CCLs (see figure 9) as well as some flux surface parameterisations useful




Figure 9. Picture-frame, Princeton-D [72] (shown here for $n_{TF} = \infty$ using a semi-analytical form, but available for finite integer n_{TF} by means of an iterative procedure), triple-arc, and poly-spline geometry parameterisations. Tangency between all wires in each parameterisation is enforced. Consistency constraints are available for each parameterisation, such that invalid combinations of variables can be eradicated during optimisation.




Figure 10. Manickam [73], Cunningham, Johner [74], and Kuiroukidis [75] flux surface geometry parameterisations.

in the design of equilibria (see figure 10), but the base class is designed to support extension to any parameterised description of geometry. Importantly, equality and inequality constraint equations can be set directly in the parameterisation and added to geometry optimisation problems.

These geometry parameterisations are such that, by dependency injection of a parameterisation, $\mathcal{S}$, and an appropriate vector of variables, α , one can obtain a primitive geometry object S with:

$$S = \mathcal{S}(\alpha, n) \quad (13)$$

where the n term denotes the number of points in the case of a parameterisation which does not have a corresponding CAD primitive (e.g. line, circle arc, B-spline), for which a B-spline is interpolated based on the number of provided points.

In general terms, the geometry object S is effectively an arbitrary curve, of the form:

$$S(t) = g(t)\hat{i} + h(t)\hat{j} + k(t)\hat{k} \quad \forall t \in [0, 1]. \quad (14)$$

Primitive wires can be discretised at any time to obtain a matrix of $\{\mathbf{x}, \mathbf{y}, \mathbf{z}\}$ coordinates, either by specific values over the range $t \in [0, 1]$ or by linear discretisation.

3.6. Materials

Materials in BLUEMIRA are handled in the separate `matproplib` package; a materials property library also developed by the authors [32]. A reasonable selection of materials with publicly available property parameterisations already exist in the library, with a wide variety of engineering, temperature-dependent thermo-mechanical and electromagnetic properties available. Additionally, a materials interface to radiation transport codes such as `OpenMC` [27] and `MCNP6` [76] is provided. Tabulation utilities for temperature-dependent material properties also exist for interfacing with FE codes.

A number of superconducting parameterisations are also available, returning the critical current as a function of field and temperature, and (occasionally) effective strand strain.

`matproplib` offers users the capability to easily create their own materials and associated property parameterisations, particularly when data or parameterisations for material properties are not publicly available. For example, we maintain a separate, private library, in which materials such as `CuCrZr`, `EUROFER97`, and `SS316-LN` are stored.

For neutronics, materials can be specified via mass fraction, atomic fraction, volume fraction, and/or as a homogenised

mixture of various materials. The nuclear cross-section data should be provided by the user. If none is provided, a default selection of cross-sections from TENDL [77] and ENDF/B-VII.1 [78], downloaded by openmc_data [79], is used. This default selection is chosen to maximise coverage over all nucleides; if a nucleide is present in both libraries, the first entry in the definition file is favoured which by default is a TENDL record.

3.7. Magnets

3.7.1. TF coil shape design. BLUEMIRA has specifically designed functionality for the common tokamak design problem of shaping the toroidal field (TF) coils in the poloidal plane. This is formulated in terms of a design optimisation problem, leveraging features in the geometry, magnetostatics, and optimisation modules. A detailed description is given in section 4.2 where the TF coil shape optimisation is presented as an example of a design procedure in BLUEMIRA.

3.7.2. Winding pack design. The magnets module provides a framework for the modelling and optimisation of superconducting magnet systems in fusion devices. While the current implementation focuses on the TF coil, future extensions are planned to include the PF coils and the CS. The methodology largely follows the approach developed within the Magnet Design Explorer (MADE) algorithm, as detailed in [80]. The framework supports both low-temperature and high-temperature superconductors, employing a hierarchical micro-to-macro modelling strategy that progresses from individual strand characterisation to the full structural integration of the coil casing.

At the millimetre-scale, individual superconducting and stabiliser strands are characterised by their density, critical current, electrical resistivity, heat capacity, and elasticity properties. These strands are assembled into cable-in-conduit conductors, where the effective electrical, thermal, and mechanical properties are homogenised. The cables are then enclosed within metallic jackets and insulating layers to form conductors. In line with the methodology adopted in the MADE framework, the anisotropic mechanical behaviour of the conductor is modelled through a spring network analogy, enabling an efficient estimation of the directional stiffness based on the homogenised response of its constituents. The mechanical stress state is subsequently evaluated by computing the local radial and axial stresses acting on the jackets and the structural case, considering both electromagnetic pressure and axial tensile forces. The Tresca yield criterion is applied for the assessment of structural integrity under electromagnetic loads.

Groups of conductors are arranged into winding packs following structured grid patterns, and their global electromagnetic and structural behaviours are derived through homogenisation. At the macroscopic scale, winding packs are integrated into the magnet structural case, where the vault, lateral plates, and closure structures are optimised to resist electromagnetic loads. This integrated framework ensures that local strand-level behaviour propagates consistently through

the cable, conductor, winding pack, and full magnet system, enabling a preliminary (0-D to 2-D) design of the superconducting magnets.

The homogenised cable structure accounts for geometrical factors such as void fraction, strand packing efficiency, and twist angle correction, which affect both thermal and electrical performance. Effective cable properties—electrical resistivity, heat capacity, and density—are computed using established mixture rules for composite materials.

The design optimisation problem solved is of the form:

$$\begin{aligned} \mathbf{u}^* &= \underset{\mathbf{u}}{\text{minimise}} : & \Delta_{\text{TF}} \\ &\text{subject to :} & \\ & & \mathbf{u}_{\min} \preceq \mathbf{u} \preceq \mathbf{u}_{\max} \\ & & T_{\text{hotspot}} \leq T_{\max} \\ & & \sigma_{\text{TF,case,peak}} \leq \sigma_{\text{TF,case,max}} \\ & & \sigma_{\text{TF,WP,peak}} \leq \sigma_{\text{TF,WP,max}} \\ & & \dots \end{aligned} \quad (15)$$

where:

$$\mathbf{u} = \{n_{\text{stab}}, \delta_{y,\text{vault}}, \delta_{\text{jacket}}\}$$

n_{stab} = number of stabiliser strands

$\delta_{y,\text{vault}}$ = TF vault lateral thickness

δ_{jacket} = TF conductor jacket thickness

Δ_{TF} = total TF thickness

T_{hotspot} = peak ‘hotspot’ temperature in the conductor

$T_{\max}$ = maximum allowable hotspot temperature in the conductor

$\sigma_{\text{case,peak|max}}$ = peak | maximum allowable Tresca stress in the TF case

$\sigma_{\text{jacket,peak|max}}$ = peak | maximum allowable Tresca stress in the TF conductor jacket

To assess thermal behaviour during fault or discharge scenarios, the cable temperature evolution is simulated through a 0-D adiabatic heat balance model:

$$\frac{dT}{dt} = \frac{J^2 \cdot \eta(T, B)}{C_p(T, B) \cdot \rho(T, B)}$$

where $J = I(t)/A_{\text{cond}}$, A_{cond} is the conducting cross-sectional area, and $B = B(t)$ denote the time-dependent current density and magnetic field. All material properties (η : electrical resistivity, C_p : specific heat capacity, ρ : density) are temperature- and field-dependent, ensuring consistency with realistic operational conditions, and are calculated using `matprop.lib`. T_{hotspot} is then computed by solving the above ordinary differential equation assuming that the current and field decay exponentially with a time-constant τ_{TFCFD} after an initial quench detection period $t_{\text{detection}}$. Figure 11 shows the behaviour of the temperature during quench subject to the posited time-dependent current and magnetic field.

At the conductor level, the cable is enclosed within a metallic jacket and an outer insulator. The effective stiffnesses in the radial and toroidal directions are computed by series/parallel homogenisation (spring network analogy) of the mechanical properties of the constituent regions. The Tresca equivalent




Figure 11. (top) conductor temperature evolution during a quench (bottom) current and magnetic field evolution during a quench, with $T_{\max} = 250$ K, $t_{\text{detection}} = 3$ s, and $\tau_{\text{TFCFD}} = 28$ s.

stress is evaluated under two primary load components: (i) radial electromagnetic pressure, $p_m = B_{\max}^2 / (2\mu_0)$, and (ii) axial force resulting from vertical electromagnetic loading.

Stress amplification factors, derived from the MADE methodology, are applied to account for local stress concentrations due to geometrical and material inhomogeneities. These factors have been validated through targeted finite element simulations (see equation (18) in [80]).

The TF case integrates multiple winding packs and structural plates (vault, lateral, inner plates). The vault radial thickness $\delta_{y,\text{vault}}$ is optimised to constrain the maximum Tresca stress at the case nose below the allowable material limit σ_{allow} , in which the hoop stress component is calculated as:

$$\sigma_{\theta}^{\text{case}} = \frac{2}{1 - \beta_{\text{vault}}^2} p_m \chi_t$$

where β_{vault} is the ratio of the outer and inner vault radii and χ_t is a stiffness-sharing factor (see equation 34 in [80]).

The optimisation problem in equation (15) is solved using a custom scheme in which the n_{stab} is first determined, satisfying the hotspot criterion, following by a second stage in which the stress constraints on the conductor jacket and case are evaluated iteratively until convergence, during which the winding pack(s) and casing are re-arranged to accommodate the updated conductor dimensions. Figure 12 shows an optimised Nb₃Sn TF coil WP and case design for a conventional aspect ratio machine.

Overall, the methodology provides a coherent, computationally efficient approach to optimise the mechanical and electromagnetic design of superconducting winding packs and coil casings.

3.8. Neutronics

BLUEMIRA has an interface to the OpenMC [27] Monte Carlo radiation transport code, which allows for two different global neutronics pipelines to be implemented: a faster axisymmetric model using constructive solid geometry (CSG), and a slower full 3-D neutronics model using faceted geometry. The former is intended for use in rapid iteration and in optimisation problems, whereas the latter is intended for the study of more detailed design points.

Both codes use the same neutron source term, see section 3.8.1, and the same material definitions, using matproplib.

At present, for a given tokamak, both pipelines return a few simple but crucial values: tritium breeding ratio (TBR), power deposition in the blanket and divertors, energy multiplication factor, and the peak damage in iron (in terms of displacements per atom per full power year, [dpa/fpy]). These parameters have important implications for the fuel cycle (see section 3.10) and power cycle (see section 3.11) analyses that can be subsequently carried out on the reactor object.

3.8.1. Neutron source. A flexible, parametric tokamak neutron source was developed by the authors, see [33], which is capable of creating a highly representative tokamak neutron source term for use in OpenMC (natively)⁹. A neutron source can be produced from arbitrary temperature and density profiles, and a plasma MHD equilibrium description (e.g. EQDSK file), without resorting to parameterisations of the plasma profiles or geometry, as is commonly done, see e.g. [81, 82]. D-T, D-D, D-³He, and T-T fusion reactions are supported.

The local source strength for a given reaction between reactants 1 and 2 is calculated as:

$$\frac{N_{12}}{dV} = \frac{1}{1 + \delta_{12}} n_1(\psi_n) n_2(\psi_n) \langle \sigma \nu \rangle_{12} (T_i(\psi_n)) \quad (16)$$

where:

N_{12} = number of reactions between reactants 1 and 2

δ_{12} = Kronecker delta (0 for different reactants, and 1 for the same)

n_1, n_2 = Ion density of reactants 1 and 2

$\langle \sigma \nu \rangle_{12}$ = velocity-averaged reactivity of reaction 1–2

The reactivities, $\langle \sigma \nu \rangle_{12}$, are calculated using ENDF/B-VIII.1 [83] cross-section data and assuming Maxwellian velocity distributions. Alternatively, for the D-T and D-D reaction channels, the Bosch–Hale fits [84, 85] can be used.

The neutron source is discretised into square cells with the normalised flux values within the LCFS being calculated directly from the equilibrium. The neutron energy spectra for the D-T and D-D reactions are modelled at given ion

⁹ Work is ongoing by contributors to the tokamak-neutron-source repository to make it available to use in MCNP6 [76] (via ‘SDEF’ card).




Figure 12. Indicative design of a conventional, Nb₃Sn TF coil WP and case, optimised using the procedure described in this section.




Figure 13. Comparison of D-T thermal fusion neutron sources with the same fusion power and temperature and density profiles, using: (a) Fausser parameterisation and (b) free boundary equilibrium. The filled contours show the total D-T reaction rates in each cell in 1/s (i.e. including the volumetric component by a factor of $2\pi x$ times the cell area), and the contour lines show the poloidal magnetic flux surfaces. The red contour represents the last closed flux surface, and the blue dot is the magnetic axis. (c) shows the absolute difference between the equilibrium and the Fausser source terms.

temperatures as per [86], whilst the T-T neutron energy spectrum is interpolated from the data presented in [87].

Figure 13 shows a comparison of D-T fusion rates for a typical conventional aspect ratio single-null tokamak between a ‘classical’ Fausser [81] parameterisation, and a free-boundary equilibrium for the same parameters and fusion power. A simple comparison of these two sources in an OpenMC model using typical tokamak geometry and typical materials showed increases of integral values dominated by interactions near the plasma (e.g. TBR) of $\approx 0.5\%$, and peak values near the plasma (e.g. peak dpa in EUROFER97) of $\approx 2\%$ when using

the equilibrium source, showing that the effects are not entirely negligible.

3.8.2. Axisymmetric global neutronics model. In this model, a simplified representation of the blanket, divertor, vacuum vessel, TF coils, and CS is constructed. This is done by representing each line segment in the true 2-D poloidal geometries as an axisymmetric CSG surface, i.e. a cone, a cylinder, a plane, or in the special case of approximating a circular arc, a torus. CSG enables faster neutron transport simulation computations than faceted geometry, but is comparatively limited in

the types of surfaces that it can represent; in particular, splines, which are often used in CAD, are not supported by CSG and must be discretised into a series of straight lines to be represented in CSG.

The axisymmetric assumption used here is acceptable for all components with the exception of the TF coil, where only the inboard leg is modelled¹⁰. Homogenised material properties are used for all solid volumes. The blankets are discretised poloidally into 3 distinct regions: first wall, breeder zone, and manifold/back supporting structure, each of which has a different material composition. An indicative CSG model of a typical tokamak geometry is shown in figure 14.

The desired volumetric integrals (e.g. TBR) are obtained by summing the responses for various cross-sections in different groups of regions.

3.8.3. Full 3-D global neutronics model. BLUEMIRA features a fully integrated pipeline capable of converting CAD into DAGMC [88] compatible models, using the `fast-ctd` code [35] also developed by the authors, which can be used by various neutronics codes.

The following steps are involved when converting CAD to DAGMC models:

- (i) **Overlap checking:** This process involves finding solids that have crossing surfaces. It begins by performing an efficient coarse-grained pair-wise search of intersecting axis-aligned bounding boxes, which are quick to generate for every component. Given the non-conformal nature of bounding boxes, the list of possible touching pairs is far greater than it should be. Further refinement is achieved by coarsely meshing each component using triangle meshes and then performing a mesh-based intersection checks using CGAL library [89].
- (ii) **Imprinting:** Imprinting is a CAD operation that creates a common surface on two solids, where they are coincident or touching. Using the list of possible touching pairs from the previous step, each pair is imprinted using OpenCASCADE.
- (iii) **Save to STP file:** The imprinted geometry is then saved to a .STP file.
- (iv) **Merging:** Merging is a CAD operation that performs a distance tolerance based identification of imprinted faces and then combines those faces into one topological entity, shared between two solids. This is required as small gaps between faces cause inaccuracies during neutronics simulations. The .step file from the previous step is ran through the merging procedure to create merged geometry.
- (v) **Surface meshing and MOAB [34]:** Every solid in the merged geometry is incrementally faceted by a 2-D Delaunay triangulation procedure in OpenCASCADE. The mesh quality and precision are controlled with two parameters: maximum linear deflection and maximum




Figure 14. Poloidal cross-section of an axisymmetric CSG model for a typical tokamak geometry. Each differently coloured region represents a different spatial sub-division and is assigned a homogenised material composition. In this example, the CS and TF are lumped together and modelled as a cylinder (magenta), the blanket (greens) is discretised into a grid of quadrilateral regions to minimise the computational complexity, and the divertor (greys) and vacuum vessel (blues) are split into 3 otherwise-polygonal regions, each with one circular arc on the plasma-facing side. The white regions are voids. A large external void region and bounding surface are included to define the simulation domain; particles crossing the outer boundary are terminated (vacuum boundary). This region is not part of the reactor design and is excluded from reported tallies.

angular deflection. During the meshing procedure, the Mesh-Oriented datABase (MOAB) library is used to tag mesh entities in the specific DAGMC format. After this is complete, a .h5m file is saved.

- (vi) **Watertight-ing:** A tool in the DAGMC package is run on the resulting .h5m model file to ensure topological entities are watertight (i.e. there are no gaps between surfaces in volumes).

Part of the work that went into building this pipeline was to establish an open-source alternative to the commercial software Cubit [90] that performs the imprinting, merging, meshing/faceting and conversion to DAGMC operations. The tools used to do this are in their own repository, `fast-ctd` [35], and are based on `overlap_checker` [91] and `occ_faceter` [92].

¹⁰ For power deposition in the TF coils, one should use a non-axisymmetric 3-D neutronics model.




Figure 15. Indicative 3-D neutronics model of a conventional aspect ratio tokamak (only surfaces are shown). The TBR per voxel (averaged over 5 batches) in the poloidal plane is overlaid (note that these voxels do not correspond with the underlying geometry used in the neutron transport simulation).

Once the geometry pre-processing is complete, the model can be run in OpenMC. Figure 15 shows a 3-D CAD model of a conventional aspect ratio tokamak, pre-processed using the pipeline described above, and the TBR resulting from a radiation transport analysis performed in OpenMC. The geometry used here is still relatively simple (e.g. the blankets are discretised into 3 regions per segment, as in the CSG model above).

It is important to note that the above workflow is best-suited for parametrically generated geometry, in which care has been taken to avoid collisions and imperfections during the parameterisation, and where small-scale detailed features (e.g. chamfers, fillets, threads, etc) are not present. In the case of manually constructed CAD models with extensive geometric detail, much of which is irrelevant to neutronics analysis, the workflow outlined above is prone to failure. Such models typically require aggressive manual de-featuring and fixing/sanitising before they can be used in the workflow.

Figure 16(a) shows the TBR results using both the CSG and CAD neutronics workflow on the same reactor design point for different numbers of neutrons, showing that reasonably converged integral results can be obtained with relatively few neutrons ($n_{\text{neutrons}} \approx 1 \times 10^5$). The pre-processing times for the CSG and CAD workflows on this design point were ≈ 10 seconds and ≈ 11 mins, respectively (not included

in the run-times shown in figure 16(b))¹¹. Note that the CAD workflow includes far more geometry, including components that do not significantly affect the result but take substantial amounts of time to pre-process (e.g. thermal shield, cryostat). Note also that the TBR calculated in the CAD neutronics model is slightly higher ($\approx 1\%$) than in the CSG model. This is not to be expected as the CSG poloidal cross-section of the blankets is volume-matched to the actual poloidal cross-section, and the 3-D CAD neutronics model includes gaps between blankets (which are not present in the axisymmetric CSG model). As the same materials and neutron source are used in both simulations, this discrepancy is most likely due to the simplified meshes presently used in the 3-D CAD neutronics workflow. Improving the meshing in the 3-D CAD neutronics workflow will be the subject of future work.

3.9. First wall heat loads

A tokamak first wall must withstand the power delivered by charged particles escaping the scrape-off layer (SOL) together

¹¹ Pre-processing timings were obtained using a single core of an 11th-generation Intel® Core™ i7-11 800 H processor [93]. The OpenMC run-times in figure 16 are given for 16 such processors on the same machine—a relatively standard modern laptop.




Figure 16. (a) Comparison of the average TBR from the CSG and CAD neutronics workflows on the same design point (5 batches). (b) Total run-times (5 batches) for the CSG and CAG OpenMC radiation transport models (i.e. excluding all geometry creation and pre-processing).

with photon radiation originating in the core and edge plasma. BLUEMIRA provides a light-weight heat load calculation procedure that, for a specified plasma equilibrium and 2-D (axisymmetric) wall contour, rapidly produces spatial maps of the incident steady-state heat flux density on the wall q_{wall} due to plasma charged particle and photon radiation. Note that the assumption of axisymmetry for the first wall is not realistic, and will tend to significantly underestimate the peak heat flux which is typically very localised. Such a procedure is, however, useful for the initial design of the poloidal contour of the first wall.

Section 3.9.1 summarises the charged particle treatment and section 3.9.2 outlines the photon radiation treatment. Full details, derivations, verification exercises, and a STEP case study are provided by Vaccaro *et al* [94].

3.9.1. Charged particle loads. The charged particle solver in BLUEMIRA is an implementation of a two-point model of the SOL, following [95, 96].

At the mid-plane the parallel heat flux density is:

$$q_{\parallel, \text{mp}}(x_{\text{mp}}) = \frac{P_{\text{SOL}}}{2\pi f_x(x_{\text{mp}})} \left(\frac{f_{\text{near}}}{\lambda_{q, \text{near}}} e^{-\frac{x_{\text{mp}}}{\lambda_{q, \text{near}}}} + \frac{1 - f_{\text{near}}}{\lambda_{q, \text{far}}} e^{-\frac{x_{\text{mp}}}{\lambda_{q, \text{far}}}} \right) \quad (17)$$

where:

x_{mp} = mid-plane radial coordinate of a flux tube

P_{SOL} = total charged particle power crossing the separatrix into the scrape-off layer (SOL), (ignoring synchrotron losses)

$\lambda_{q, \text{near}}, \lambda_{q, \text{far}}$ = near and far SOL power decay length constants

f_{near} = fraction of SOL power assumed to be deposited in the near SOL

$f_x(x_{\text{mp}})$ = local total flux tube expansion factor

the SOL power is divided up by user-specified fractions to the inboard and outboard SOL legs, and in the case of a double-null, upper and lower fractions.

Each flux surface is followed from the mid-plane to its first intersection with the wall contour; any later crossings are ignored. This single-hit rule reproduces geometric shadowing and ensures exact power accounting, because every flux tube deposits its energy once and only once, so the sum of wall loads equals the input P_{SOL} .

The heat flux density delivered to the wall by charged particles is:

$$q_{\text{wall}, c} = q_{\parallel, \text{mp}} \frac{B_{p, \text{mp}}}{B_{p_i}} \sin(\delta_i) \quad (18)$$

with B_p the poloidal field and δ_i the local poloidal incidence angle.

The approach only explicitly computes the parallel component, $q_{\parallel, \text{mp}}$; cross-field transport is accounted for by the decay lengths λ_q . Consequently, the last SOL tube originates outside the physical wall and its parallel heat cannot be deposited. Many tracing codes combat this with heavy mesh refinement or by smearing the ‘missing’ power over all tubes. BLUEMIRA instead transfers that power perpendicularly to the wetted zone around the mid-plane where cross-field diffusion is most relevant restoring energy balance without numerical smoothing and producing realistic mid-plane peaks for the first wall design (figure 17). The implementation has been benchmarked against the 3-D field-line tracer SMARDDA [97, 98] for an axisymmetric EU-DEMO 2017 case, showing close agreement in both intersection pattern and heat flux profile while conserving total power.

3.9.2. Radiation loads. Photon heating is treated in two stages: (i) a volumetric radiative power map inside the closed and open field line regions is built and (ii) photon transport to the first wall surface is calculated with the Monte-Carlo ray-tracer CHERAB [99, 100].

Core plasma. For a user-specified set of electron temperature and electron density profiles, and impurity fractions (constant concentration across the profile), the local radiated power is:

$$\frac{dP_{\text{rad}}}{dV} = \sum_k f_k n_e^2 L_k(T_i) \quad (19)$$

where $f_k = n_k/n_e$ is the fractional abundance of impurity ion k and L_k is the tabulated power loss function for said ion drawn from ADAS [101], which includes line radiation from heavy species and Bremsstrahlung. Synchrotron radiation is not modelled here.

SOL and divertor. A two-point model [95, 96] is solved along the separatrix to obtain electron temperature and density at four key locations:

- (i) upstream (mid-plane),
- (ii) onset of the radiation zone (just above the X-point),




Figure 17. A conventional aspect ratio double null equilibrium showing: (a) Charged particle heat loads on the main chamber first wall. The red lines show the open flux surfaces carrying charged particles. (b) Core radiation source term and heat flux density on the first wall, as calculated by CHERAB. (c) SOL radiation source term and heat flux density on the first wall, as calculated by CHERAB.

- (iii) entrance to the recycling zone (just below the X-point),
- (iv) the divertor target.

target temperature, or detachment threshold are exposed to the user.

The mid-plane temperature is computed for the conduction-limited case and is the only location that requires an explicit density input. At each subsequent location the density is inferred from the previous one by enforcing constant total pressure, i.e. $n_{e,i} = n_{e,i-1} T_{e,i-1} / T_{e,i}$. When the target temperature falls below 5 eV, a fractional momentum-loss factor f_m is introduced to account for ion-neutral friction, modifying the local pressure balance and mimicking fully detached solutions.

Once the separatrix temperature and density are determined, the corresponding values on all SOL flux surfaces are obtained by applying radial exponential decay. Power dispersal follows: $q_{\parallel}(x) = q_{\parallel,0} \exp(-x/\lambda_q)$, where the parameter λ_q controls the radial power spreading. Analogous exponential profiles are used for $T_i(x)$ and $n_i(x)$ with their respective user-input decay lengths.

Between each of the four key locations (mid-plane to target) the electron temperature falls and the density rises according to exponential or Gaussian laws (options supplied to the user), while pressure balance and the upstream forcing function ensure consistency with detachment physics when $T_{\text{target}} \lesssim 2\text{ eV}$. Access to tabulated impurity power loss rates from an atomic database (e.g. ADAS) allows the code to assemble a full 2-D map of radiated power density MW m^{-3} ; the resulting radiation typically concentrates near the X-points, see figure 17(c). Parameters such as the radiation zone length,

Photon transport. The combined core-plus-SOL source is discretised on an axisymmetric grid and passed to CHERAB, which is used to calculate the incident surface power density $q_{\text{wall},\gamma}$ (figure 17). Because the radiated power calculation is analytic/table-driven, and the ray-tracing exploits variance-reduction techniques, the full radiation map is obtained quickly enough for iterative reactor design.

3.10. Fuel cycle

The fuel cycle module in BLUEMIRA is described in [102]. It is a simplified dynamic fuel cycle model that, for a given set of inputs and operational timelines, calculates the tritium start-up inventory, $m_{T,\text{start}}$, the inventory inflection point, t_{infl} , and the reactor doubling time, t_d . The fuel cycle model is based on a three-loop direct internal recycling (DIR) architecture, as described in [103]. Many of the inputs are ‘high-level’ (e.g. separation fractions, residence times) and should be obtained from detailed fuel cycle component design and modelling.

At present, the architecture of the modelled fuel cycle is not very configurable. That said, the relative simplicity of the model, and the wide variety of input parameters, enables comparisons of a relatively broad set of parameter and design space. Furthermore, any of the three loops can quite easily be switched off by setting their flow fractions to zero, thus to




Figure 18. (top) Total tritium inventories in the reactor over the reactor lifetime, with the inventory inflection point and reactor doubling time annotated; (bottom) Sequestered tritium inventories in the in-vessel environment (incl. first wall), breeding blanket (excl. first wall), and tritium fuelling and vacuum (TFV) systems.

some extent capturing different fuel cycle configurations, see [102].

Crucially, the fuel cycle model includes simplified, configurable models of tritium retention which depend on the incident flux of tritium and some saturation limit. These are used for the absorption of tritium in the first wall, and also to model tritium retention inside various components in the tritium plant. These sinks of tritium play a major role in determining the start-up inventory. Since [102], an additional sink model was included with a square-root dependence on inventory, based on some more detailed simulations carried out in [104], see [4] appendix D. Note that the inputs to such models should always be derived in some form from higher fidelity tritium trapping and transport simulations, which take into account a variety of design, material, and operational considerations (e.g. temperature, neutron damage, etc). This is in contrast to a more recent fuel cycle model [105] which accounts for trapping directly without use of surrogates.

The fuel cycle can be run with arbitrary operational timelines to support the more accurate estimation of the start-up inventory in first-of-a-kind machines (characterised by low load factors). Typically the inventory inflection point occurs early on in the operation of a machine, and depends heavily on the duty cycle of the machine. In particular, lengthy unplanned shutdowns in the early phases of operation can have a significant impact on the start-up inventory. Functionality is provided to create pseudo-randomised operational timelines following general load factor trends. The fuel cycle model can be run in Monte Carlo fashion, generating distributions for the parameters of interest.

Figure 18 shows an indicative result for a typical DIR fuel cycle in an EU-DEMO-like machine, with a progressively increasing load factor (0.1 to 0.4 over the plant lifetime).

This fuel cycle model is intended primarily as a loosely-coupled post-processing step, which can be performed on a reactor, or separately, as a standalone analysis.

3.11. Power balance

A simple 0-D fusion reactor steady-state (during flat-top) power balance model is also available in the codebase. Figure 19 shows an indicative result for this model. Much like the fuel cycle model, it is relatively specific a specific architecture of fusion reactor power cycle, and expects a variety of input parameters from a reactor object. This simplified model is a crude but useful addition; it is slightly more detailed than the PROCESS power balance, and can be passed the results from higher fidelity calculations (e.g. neutronics analyses, see section 3.8.2). Most of the parasitic loads for the various plant sub-systems are taken from PROCESS, in the absence of better simplified models. A complete description can be found in [4], appendix C.

4. Design procedures in BLUEMIRA

In this section, we discuss how design procedures are implemented in BLUEMIRA.

4.1. Optimisation problems

Optimisation is used frequently throughout the various design procedures in BLUEMIRA. As such, it is useful to have a generic approach to setting up an optimisation problem of the form:

$$\begin{aligned}
 \mathbf{u}^* &= \underset{\mathbf{u}}{\text{minimise}} : & f(\mathbf{u}) \\
 &\text{subject to} : & \\
 & & \mathbf{u}_{\min} \preceq \mathbf{u} \preceq \mathbf{u}_{\max} \\
 & & g(\mathbf{u}) \preceq \mathbf{0} \\
 & & h(\mathbf{u}) = \mathbf{0} \\
 & & \dots
 \end{aligned} \tag{20}$$

We define $\mathbf{u}$ as a generic vector of design variables, and use the superscript $*$ to indicate optimality.

By the ‘no free lunch’ theorem [106], no particular optimisation algorithm will out-perform all others over all problems. Some algorithms are better suited to certain problems. In BLUEMIRA, an abstract interface is used for single-objective optimisation problems, such that any optimisation package can be used. At present, two optimisation suites, NLOpt [107] and scipy [108] are wrapped by this abstract interface, providing good selection of gradient-based and stochastic algorithms are available, many of which can handle the above problem formulation (i.e. with equality and inequality constraints). NLOpt algorithms are more frequently used, as we have found that this optimisation library is the most extensive (in terms of the algorithms available) and robust (in terms of implementation).

Objective and constraint functions can be used with analytical gradients, provided they are known, in order to reduce the number of function evaluations. If a gradient-based optimisation algorithm is used, but an analytical gradient is not provided, a numerical gradient is calculated automatically using central finite difference. This is not always a robust approach, but for some problems it can be helpful.




Figure 19. Indicative steady-state power balance model result for a typical tokamak with a helium-cooled ceramic breeder blanket.

4.2. Example of a design procedure: TF coil design

As a framework in which the user can combine a variety of multi-disciplinary activities to produce a reactor design, BLUEMIRA exposes a handful of key software constructs to the user:

- (i) Parameter, ParameterFrame
- (ii) Designer
- (iii) Builder
- (iv) Component
- (v) Reactor

Figure 20 shows these constructs arranged in a typical workflow. These will be introduced by means of a concrete example: that of the design of a TF coil.

A Parameter holds a particular value, with its associated name (in long and short forms, for text and code respectively), description, unit, and source. A Parameter may be updated and its values and source overwritten, which is subsequently tracked in its history. This is useful when a particular value of interest has some initial estimate, which is subsequently overwritten by a more detailed calculation.

ParameterFrames are sets of Parameters (with a few convenient features) that are used as explicit interfaces between calculation procedures. Here for the design and build of the TF coil, a ParameterFrame containing values relating to the geometrical build of the TF coil, winding pack, insulation thickness, etc as well as some machine design parameters such as R_0 , B_T , and the maximum allowable ripple in the plasma, $\delta_{\phi, \max}$, are used as inputs to a Designer.




Figure 20. Abstraction of a typical BLUEMIRA workflow producing a Component as part of a Reactor data structure.

A Designer is responsible for producing a design given some inputs, and is typically some form of design optimisation problem. Here the optimisation of the centreline shape of a TF coil winding pack is carried out by a Designer, which takes the aforementioned ParameterFrame as an input, along with a shape parameterisation for the TF CCL, and two geometrical objects: the plasma LCFS and a keep-out zone (KOZ) which the TF CCL must be outside of. It returns the optimised TF coil CCL.




Figure 21. Diagram of the TF coil CCL design optimisation problem, showing the initial and optimal CCL Triple-Arc shapes. The TF ripple is shown for the optimal CCL.

The approach for the TF coil shape design taken here is to minimise the length of the CCL, subject to a KOZ constraint, and such that the ripple on the plasma LCFS is below $\delta_{\phi, \max}$, see equation (21) and figure 21. A ‘triple-arc’ up-down symmetric shape parameterisation, similar to the ITER TF coil shape, is used such that minimum bending radii can also be set.

Of course, other objectives (e.g. minimisation of bending) and constraints (e.g. toroidal displacement under loads) are possible. At present, this approach makes no consideration of in-plane bending in the TF coil structure as the subsequent inclusion of the PF coils introduces significant local bending into the TF coil shape, negating to some extent efforts to reduce in-plane bending due to TF loads alone. Nevertheless, a ‘bending-free’ shape parameterisation for the TF CCL can be used to solve the above optimisation problem, without introducing any further constraints.

$$\begin{aligned} \alpha^* = \underset{\alpha}{\text{minimise}} : & \int_0^1 \sqrt{\left(\frac{dx}{dt}\right)^2 + \left(\frac{dz}{dt}\right)^2} dt \\ \text{subject to :} & \\ & \{r_1, r_2, r_3\} \succcurlyeq r_{\text{TF}, \min} \mathbf{1} \\ & \theta_1 + \theta_2 + \theta_3 \leq 180^\circ \\ & \delta_\phi|_{\partial\Omega_p} \preccurlyeq \delta_{\phi, \max} \mathbf{1} \\ & d^\pm(S_{\text{CCL}}(\alpha, n), \Omega_{\text{KOZ}}) \preccurlyeq \mathbf{0} \end{aligned} \quad (21)$$

where:

$\alpha = \{r_1, r_2, r_3, \theta_1, \theta_2, \theta_3, l_{\text{straight}}\}$ are the radii and angles of the three circular arcs, and l_{straight} is the length of the straight inboard leg

$r_{\text{TF}, \min}$ = minimum bending radius

$\delta_\phi|_{\partial\Omega_p}$ is the TF ripple on the LCFS, $\partial\Omega_p$, calculated as per equation (22)

$d^\pm(S(\alpha, n), \Omega_K)$ is the signed distance function from the shape S to the exclusion zone Ω_K , calculated as per equation (24)

For each function evaluation, the CCL is created from the shape parameterisation, which is used to create a magnetostatic model of the TF cage, using 3-D trapezoidal prism current sources to represent the winding pack (see section 3.3). The toroidal field ripple, δ_ϕ , is then calculated as:

$$\delta_\phi(x, z) = \frac{B_t(x, \phi = 0, z) - B_t(x, \phi = \pi/n_{\text{TF}}, z)}{B_t(x, \phi = 0, z) + B_t(x, \phi = \pi/n_{\text{TF}}, z)} \quad (22)$$

The peak TF ripple can be formally evaluated as the maximum ripple on the LCFS, see equation (23), but a number of alternative heuristic approaches are available for parametrically specifying the point(s) on the LCFS at which to evaluate the ripple¹²,

$$\delta_{\phi, \text{peak}} = \underset{t}{\text{maximise}} : \delta_\phi|_{\partial\Omega_p}(t). \quad (23)$$

The KOZ constraint is implemented in the form of a discretised signed distance function (see equation (24)) from the TF CCL, S_{CCL} , to some region, Ω_{KOZ} . Here, the KOZ is a region, offset from the vacuum vessel thermal shield, and includes the TF coil thickness from the plasma-facing casing to the CCL.

$$\begin{aligned} d^\pm(S, \Omega_{\text{KOZ}}) &= d^\pm(x, z, \Omega_{\text{KOZ}}) \quad \forall x, z \in S \\ d^\pm(x, z, \Omega_{\text{KOZ}}) &= \min s \sqrt{(x - x_\partial)^2 + (z - z_\partial)^2} \quad \forall x_\partial, z_\partial \in \partial\Omega_{\text{KOZ}} \\ s &= 1 \text{ if } x, z \in \Omega_{\text{KOZ}} \text{ else } -1. \end{aligned} \quad (24)$$

Following on from the Designer, a Builder is responsible for constructing a CAD model given some inputs. The TF coil CAD model Builder takes as inputs the ParameterFrame mentioned above, and the TF CCL optimised by the Designer. It builds the CAD models for the TF WP, ground insulation, and casing. Here a simple wedged case design is used, with an ample space reservation in the inboard case corners (for detailed features such as pre-compression rings, poloidal shear keys, etc), see figure 22. Inter-coil structures and PF coil supports are added in a separate procedure.

The Builder produces a Component as an abstraction of real-world system, which may or may not have a physical embodiment. Components are organised in a classical tree structure. Components that represent physical objects (e.g. TF casing), have an associated CAD model and material. The granularity of allocations of CAD models and materials becomes important in subsequent finite element or radiation

¹² Technically, the minimisation of an objective with a constraint on the peak ripple (maximum of the ripple on the LCFS) is the correct formulation of this problem, but can be considerably accelerated if one knows more about the problem. For example, with picture frame TF coils and a positive triangularity LCFS, the peak ripple will invariably be located at the outboard midplane of the LCFS, and there is generally no need to evaluate the ripple elsewhere.




Figure 22. CAD model of a single TF coil produced by the Builder, showing the casing as transparent and the winding pack. The ground insulation is not shown for clarity.

transport analyses. It is up to the user whether or not to use smeared properties for a given volume; a winding pack, for example, which can be further broken down into ground insulation, turns, etc.

All the Components for the various modelled sub-systems are passed to the root of the tree, the Reactor, which is the complete data structure for a design point.

5. Case study: parameterising a conventional aspect ratio reactor design in BLUEMIRA

In this section we briefly describe a full reactor design parameterisation in BLUEMIRA, taking as a worked example a conventional aspect ratio tokamak, similar in nature to an EU-DEMO design, see e.g. [8, 109].

Parameterising a reactor design bears a significant up-front cost, in that one must specify the desired sequence of activities, the formulation of the various design optimisation steps, and the parameterisation of the geometry of various sub-systems. Once done, however, it enables a variety of design exploration studies to be carried out relatively effortlessly.

The design parameterisation presented here produces a fairly complete description of a tokamak with ‘intermediate’ fidelity results, and loosely replicates the design procedure performed by the EUROfusion programme management unit and its associates when generating a ‘baseline’ design point (which takes several months). Indeed, a BLUEMIRA design parameterisation similar to the one presented here was used to generate an initial design point for the proposed

‘low-aspect ratio’ revision to the EU-DEMO design [20], using the PROCESS design point described in [110].

For the purposes of this demonstration, we have used the following sequential design workflow:

- (i) PROCESS run
- (ii) Initial ‘reference’ equilibrium design
- (iii) Initial FW shape optimisation
- (iv) TF coil poloidal shape optimisation
- (v) PF system design
- (vi) Vertical stability assessment
- (vii) FW shape optimisation and panelling
- (viii) CAD component design (blankets, divertors, vessel, etc)
- (ix) Axisymmetric CSG neutronics analysis
- (x) CAD component design (PF coils, ports, thermal shields, etc)
- (xi) Steady-state power balance model

Clearly, some aspects differ from the original procedure used in EUROfusion. In part, this is due to the different codes used in the baseline design procedure, and them not necessarily being available in BLUEMIRA or that the alternatives in BLUEMIRA are not as good or as fully featured. However, many of the design procedures for the EUROfusion baseline are done manually and by different people, not all of whom are able to describe their work in terms of equations or algorithms. In particular, the design of some of the less-critical geometry (e.g. poloidal cross-section of the vacuum vessel, PF supports, etc) where design objectives and constraints are less

design-driving, it comes down to choice of shape parameterisations and various other subjective decisions.

5.1. 0/1-D systems code optimisation with PROCESS

Here, we follow the typical way in which PROCESS used to be run for EU-DEMO designs (see e.g. [109]), namely the minimisation of the major radius, subject to a variety of constraints and bounds, see equation (25). This differs substantially from how PROCESS has been used most recently on EU-DEMO designs, see [110].

$$\begin{aligned}
 \mathbf{u}^* &= \underset{\mathbf{u}}{\text{minimise}} : R_0 \\
 &\text{subject to :} \\
 &\mathbf{u} \succcurlyeq \mathbf{u}_{\min} \\
 &\mathbf{u} \preccurlyeq \mathbf{u}_{\max} \\
 &P_{\text{el,net}} = 500 \text{ MWe} \\
 &\tau_{\text{flatop}} \geq 2 \text{ h} \\
 &P_{\text{sep}} B_T / q_{95} A R_0 \leq 9.2 \text{ MW.T/m} \\
 &1.1 \leq f_{\text{LH}} \leq 1.2 \\
 &\bar{n}_e \leq 1.2 n_{\text{GW}} \\
 &n_{\text{e,core}} \geq n_{\text{e,ped}} \\
 &\beta \leq \beta_{\max} \\
 &\sigma_{\text{TF,case,peak}} \leq 580 \text{ MPa} \\
 &\sigma_{\text{TF,WP,peak}} \leq 580 \text{ MPa} \\
 &J_{\text{TF}} \leq J_{\text{TF,max}} \\
 &\Delta_{\text{T,TF}} \geq 1.5 \text{ K} \\
 &\sigma_{\text{CS,peak}} \leq 660 \text{ MPa} \\
 &\Delta_{\text{T,CS}} \geq 1.5 \text{ K} \\
 &\sigma_{\text{VV,TFCFD}} \leq 93 \text{ MPa} \\
 &\dots
 \end{aligned} \tag{25}$$

where the variable vector is:

$$\begin{aligned}
 \mathbf{u} &= \{\langle T_e \rangle, \langle n_e \rangle, q_{95}, B_T, \Delta_{\text{bore}}, \Delta_{\text{CS}}, \Delta_{\text{CS} \rightarrow \text{TF}}, \Delta_{\text{TF}}, \Delta_{\text{TF,nose}}, \\
 &\delta_{\text{TF,turn}}, I_{\text{TF,turn}}, \tau_{\text{TF}}, J_{\text{CS,EOF}}, f_{\text{Xe}}, \dots\} \\
 \mathbf{u}_{\min} &= \text{lower bounds} \\
 \mathbf{u}_{\max} &= \text{upper bounds} \\
 q_{95} &= \text{safety factor at the 95th percentile flux surface} \\
 \Delta_{\text{bore}} &= \text{bore thickness} \\
 \Delta_{\text{CS}} &= \text{CS thickness} \\
 \Delta_{\text{CS} \rightarrow \text{TF}} &= \text{gap thickness between CS and TF} \\
 \Delta_{\text{TF,nose}} &= \text{Inboard TF nose thickness} \\
 \delta_{\text{TF,turn}} &= \text{TF turn width} \\
 I_{\text{TF,turn}} &= \text{TF turn current} \\
 \tau_{\text{TF}} &= \text{TF discharge time constant} \\
 J_{\text{CS,EOF}} &= \text{CS current density at EOF} \\
 f_{\text{Xe}} &= \text{Xenon impurity fraction}
 \end{aligned}$$

Consistency equality constraints for the radial build, plasma beta, and plasma power balance are also used, but not shown in equation (25) as they relate to the self-consistency of the solution rather than the design intent.

and the terms in the constraints are:

$$\begin{aligned}
 P_{\text{sep}} &= \text{power crossing the separatrix} \\
 f_{\text{LH}} &= P_{\text{sep}} / P_{\text{LH}} \text{ where } P_{\text{LH}} = \text{is the L-H transition power} \\
 \bar{n}_e &= \text{line-averaged electron density} \\
 f_{\text{GW}} &= \text{Greenwald fraction} \\
 n_{\text{GW}} &= \text{Greenwald density} \\
 n_{\text{e,core}} &= \text{core electron density} \\
 n_{\text{e,ped}} &= \text{pedestal electron density} \\
 \Delta_{\text{T,TF}} &= \text{temperature margin in the TF superconductor} \\
 \Delta_{\text{T,CS}} &= \text{temperature margin in the CS superconductor} \\
 \sigma_{\text{VV,TFCFD}} &= \text{Hoop stress in the vacuum vessel during a TF coil fast discharge}
 \end{aligned}$$

The solution of the PROCESS optimisation problem (i.e. the variable vector, radial build, profiles, etc) are used as inputs in the subsequent BLUEMIRA design stages.

5.2. Initial reference equilibrium and FW design

An initial ‘reference’ free-boundary equilibrium is required in order preliminarily position the FW, blankets, and vacuum vessel.

This is created by setting up a series of constraints on the plasma position (isoflux, field null, etc) and solving with some preliminary PF coil positions in order to match the description of the equilibrium from the 0-D parameters from PROCESS (R_0 , A , κ , δ). A Kuroikidis [75] single-null shape parameterisation for the desired LCFS is used.

A preliminary FW design is then created using a purely geometric set of rules (i.e. without any heat flux calculations). Here the FW is defined simply as the convex hull of a flux offset from the LCFS, and a geometric offset from the LCFS. A classical ITER-like divertor shape parameterisation is for the divertor plasma-facing surfaces.

5.3. Toroidal field coil shape optimisation

See section 4.2.

5.4. Poloidal field system optimisation

In this reactor parameterisation, the PF coil centrepoinets are positioned along a ‘track’, offset from the outer edge of the TF coil casing, in the poloidal plane. Some KOZs, inside of which the PF coils cannot be placed, are defined, and the coil positions are optimised using a two-stage approach, in which the positions and the currents of the PF coils are optimised separately.

5.4.1. Remote maintenance KOZs. Here we assume a vertical maintenance architecture, see [111, 112], where it is vital that the upper port is large enough to extract the in-vessel blanket segments. In equation (26), a constrained optimisation problem is solved in which the size of the upper port is minimised, with some simplified constraints ensuring that the inboard and outboard blanket segments can be extracted from the upper port (with no rotations of blanket segments about the y-axis).

$$\begin{aligned}
\mathbf{u}^* &= \underset{\mathbf{u}}{\text{minimise}} : & x_{\text{ob}} - x_{\text{ib}} \\
&\text{subject to :} & \\
& x_{\text{ob}} - x_{\text{cutBSS}} - \Delta_{\text{RM}} \leq x_{\text{BB}_{\text{max}}} - x_{\text{cutBSS}} \\
& x_{\text{cutFW}} - x_{\text{BB}_{\text{max}}} \leq x_{\text{ob}} - x_{\text{cutFW}} + \Delta_{\text{RM}} \\
& x_{\text{ib}} + \frac{|x_{\text{cutFW}} - x_{\text{BB}_{\text{min}}}|}{2} \leq x_{\text{cutFW}} \\
& \gamma_{\text{min}} \leq \gamma_{\text{cut}} \leq \gamma_{\text{max}}
\end{aligned} \tag{26}$$

where:

$\mathbf{u} = \{x_{\text{ib}}, x_{\text{ob}}, x_{\text{cutFW}}, \gamma_{\text{cut}}\}$
 x_{ib} = minimum inboard upper port radius
 x_{ob} = maximum outboard upper port radius
 x_{cutFW} = minimum radius of the cut location in the BB FW
 x_{cutBSS} = minimum radius of the cut location in the BB back supporting structure
 γ_{cut} = cut angle between the inboard and outboard BB segments
 $x_{\text{BB}_{\text{max}}}$ = maximum radius of the outboard BB
 $x_{\text{BB}_{\text{min}}}$ = minimum radius of the inboard BB
 Δ_{RM} = remote maintenance clearance around BB segments during maintenance

The equatorial port opening size is taken as an input (typically 2 m by 2 m), and the lower port is sized based on the divertor size and a user-input angle (typically -15° to -45°).

Geometric offsets to allow for the thickness of the vacuum vessel port walls, thermal shield, and various gaps are included to produce the KOZs for the PF coils.

5.4.2. PF coil position optimisation problem. The KOZs for the ports, described in the previous section, and an additional KOZ for the TF coil gravity supports are used to define sub-segments of the ‘track’ along the TF coil in which PF coils can be placed, as described in [57]. The parameterisation positions of the coils, $\mathbf{L}$, are optimised such that:

$$\begin{aligned}
\mathbf{L}^* &= \underset{\mathbf{L}}{\text{minimise}} : & \|\mathbf{I}^*\|_2^2 \\
&\text{subject to :} & \\
& \mathbf{L}_{\text{min}} \preceq \mathbf{L} \preceq \mathbf{L}_{\text{max}}
\end{aligned} \tag{27}$$

where $\mathbf{I}^*$ is the vector of optimal currents for a given set of positions, $\mathbf{L}$, which is the result of a current optimisation problem, discussed in the next section.

5.4.3. PF coil current optimisation problem. For the PF coil current optimisation, a similar optimisation problem to the one described by Albanese *et al* in [113] for EU-DEMO is solved, as detailed in [57], see equation (28).

$$\begin{aligned}
\mathbf{I}^* &= \underset{\mathbf{I}}{\text{minimise}} : & \|\mathbf{GI} - \mathbf{b}_t + \mathbf{b}_p\|_2^2 + \|\mathbf{I}\mathbf{I}\|_2^2 \\
&\text{subject to :} & \\
& |\mathbf{I}| \preceq \mathbf{I}_{\text{max}} \\
& B_p|_{x_{\text{point}}} \leq 0 + \epsilon \\
& \mathbf{B}_{p,\text{coils}} \preceq \mathbf{B}_{p,\text{max}} \\
& |\mathbf{F}_{z,\text{coils}}| \preceq \mathbf{F}_{z,\text{max}}
\end{aligned} \tag{28}$$

This optimisation problem is solved for the required plasma boundary flux at start of flat-top (SOF) and end of flat-top (EOF), for a desired pulse length. This approach differs from the approach in [113] in two ways: (i) the objective function is a minimisation of error (regularised with coil currents), rather than a minimisation of the coil currents and a constraint on the allowable error, and (ii) the EOF boundary flux is an input rather than being the result of a minimisation of the flux at EOF.

The CREATE approach is a more formal and precise method using functionality in the CREATE-NL+ code [114], whereas the approach described here is capable in most cases of providing a useful result, unsupervised. Work is ongoing to improve the approach presently used in BLUEMIRA, such that the pulse length can be calculated subject to an inequality constraint on the geometrical error with respect to a prescribed plasma shape.

5.5. Vertical stability assessment

The vertical stability calculation described in section 3.4.3 is performed for a representative equilibrium, including the passive contributions from the vacuum vessel, returning an estimate of m_s . In this design parameterisation, the vacuum vessel is a toroidally continuous double-shelled structure (with a shell thickness of 60 mm), through which ports are placed. Here, we conservatively assume that only the outer VV shell contributes to the passive vertical stabilisation, however we assume that electrical connections are placed in each port at the position where the outer VV shell’s toroidal continuity would otherwise be broken. In reality, of course, although this could be achieved with specially designed attachments in the port shield plugs, such conducting paths would not be so uni-formally placed or as effective. Furthermore, the contributions of the active, controlled (PF and CS) coils are not included in the stabilising term calculated here.

5.6. First wall shape

With equilibria for the start and end of flat-top, the first wall is re-designed. The convex hull of some geometric and flux offsets from both SOF and EOF LCFSs is taken, and is split into inboard and outboard segments at a radial position calculated from the remote maintenance KOZs. The inboard and outboard first wall poloidal ‘silhouettes’ are then ‘panelled’ into flat surfaces, see equation (29).

$$\begin{aligned}
\mathbf{t}^* = \underset{\mathbf{t}}{\text{minimise}} : & \sum_{i=0}^{n_{\text{panels}}} L_{p_i} \\
\text{subject to :} & \\
& \alpha_p \preccurlyeq \alpha_{\max} \mathbf{1} \\
& L_p \succcurlyeq L_{\min} \mathbf{1}
\end{aligned} \tag{29}$$

where:

$\mathbf{t}$ = is the vector

$\in [0, 1]$ = of parameterised joint locations between the panels

n_{panels} = is the number of panels

α_p = is the vector of angles between the panels

$\alpha_{\max}$ = is the maximum angle between panels

L_p = is the vector of panel poloidal lengths

$L_{\min}$ = is the minimum panel poloidal length

The panels are parameterised to sit tangent to the initial convex hull, thus doing away with the need for an explicit keep-out-zone constraint. If the above optimisation problem is not found to have a satisfactory solution for a given n_{panel} , n_{panel} is incrementally increased until the constraints can be met.

The divertor plasma-facing surfaces are crudely defined for some specified inboard and outboard leg lengths and grazing angles, along with some normalised flux offset for the private flux region mimicking a closed divertor geometry. More refined approaches to defining the divertor geometry are possible.

5.7. Global neutronics

The axisymmetric CSG neutronics model (see section 3.8.2) is used in this procedure to improve on the PROCESS estimates of the power deposition in the in-vessel components, and energy multiplication, with the aim of providing a better estimate of the global thermal power balance in the reactor. The blankets are radially discretised into three regions: first wall, breeding zone, and manifold and back-supporting structure, each with their own homogenised material composition. Typical thicknesses and material compositions for these regions are taken from detailed helium-cooled ceramic breeder (HCCB) blanket studies [115, 116]. Two batches of 100 000 neutrons were run for each point.

5.8. Power balance

A simplified, 0-D, steady-state power balance model is available in BLUEMIRA, at a similar level of definition and depth to the one found in PROCESS. Indeed, some of the estimates for the parasitic loads in particular are simply taken from PROCESS. However, with results from actual neutronics analyses for energy multiplication and power deposition in the IVCs, we are able to provide slightly higher fidelity in terms of the overall reactor power balance. The PROCESS values for the parasitic loads for e.g. magnet supplies, the cryoplant, the tritium plant, etc are used.

6. Case study: performing design exploration study in BLUEMIRA

Using the conventional tokamak design parameterisation and procedure outlined in the previous section, we perform a scan in aspect ratio ($A \in [2.6 - 3.6]$) for varying number of TF coils ($n_{\text{TF}} \in \{15, 16, 17, 18\}$) to demonstrate the ability to rapidly carry out higher-fidelity design studies.

6.1. Results

As a comparison with a lower-fidelity study, we show the PROCESS code results (which are used as the starting point for the BLUEMIRA study), see figure 23. For this relatively simple design problem in PROCESS, a slight minimum in the minimisation objective function, R_0 , can be seen here between $A = 3.0$ and 3.1 , depending on the number of TF coils.

Figure 24 shows the reactor geometry of three design points in the scan, where the variation in plasma geometry is clearly visible. At lower A , the toroidal field on-axis is lower for the same fusion power, leading to a thinner TF coil, whereas at higher A the toroidal field is higher, leading to thicker TF coils.

6.1.1. Comparison between BLUEMIRA and lower-fidelity PROCESS models. In this section, we draw a number of comparisons between some lower-fidelity models used in PROCESS with some higher-fidelity formulations BLUEMIRA, highlighting some inaccuracies and limitations in classical systems codes, whilst commenting on the potential implications in terms of the design space.

We compare the net electric power, $P_{\text{el,net}}$, as computed using the initial PROCESS results (in which a constraint of $P_{\text{el,net}} = 500$ MWe is used), and the subsequent BLUEMIRA results, see figure 25.

Here, in BLUEMIRA, we calculate $P_{\text{el,net}}$ in a similar way to PROCESS, but using the results from the CSG neutronics workflow to determine power deposition in the in-vessel components. The results are only subtly different, showing that the simplified neutron power multiplication and distribution model in PROCESS is reasonably accurate. The energy multiplication factor calculated using the CSG neutronics model was not found to vary significantly with aspect ratio, and averaged 1.347, comparing well with the PROCESS input of 1.350 used here. The differences are largely down to different treatment of how the divertor geometry is considered. The parameterisation for the divertor used in BLUEMIRA here is a classical ‘ITER-like’ cassette (see figure 24), which takes up proportionally more space at higher aspect ratio and leads to less power being deposited in the blankets (which are cooled at higher temperatures). PROCESS calculates the neutron power deposited in the divertor as a constant fraction of the total neutron and core radiated power, which is a user input and does not change with A . This explains the different trends seen in figure 25. This information could relatively easily be used to improve the model presently in PROCESS.

The fluctuations in the BLUEMIRA results primarily arise from the geometry discretisation in the CSG neutronics model




Figure 23. PROCESS minimisation objective values for the $\{A, n_{TF}\}$ design space explored. The red crosses indicate failure in the subsequent BLUEMIRA runs.

and the statistical uncertainty associated with the limited number of neutron histories and batches. Here the number of TF coils has little effect on the results, and we only show results for $n_{TF} = 16$, for clarity.

The peak field on the TF coil as calculated using a 3-D magnetostatic model (including 3-D and PF contributions) was found to be consistently higher than the peak field reported in PROCESS, calculated as $1/r$, see figure 26. PROCESS does have a surrogate model for the peak field including 3-D effects. Figure 26 shows this estimate is actually more reasonable, and in these cases (for the ‘triple-arc’ TF coil shape) overestimates the peak field slightly. This has implications in terms of the achievable critical current density in the superconductor, the quench behaviour of the TF WP, and the radial build as a whole.

The peak TF ripple on the separatrix is generally poorly estimated by PROCESS, particularly at higher aspect ratio, as can be seen from figure 27.

PROCESS has a simplified TF coil shape parameterisation with a single variable (the outboard limb radius) which is calculated based on a fit to some magnetostatic results to achieve the specified level of ripple at the outboard mid-plane (in this case 0.6%). Clearly, however, this simplified model does not accurately reflect the true peak TF ripple along the LCFS in this parameterisation. The ripple calculation using a series of rectangular cross-section current sources to represent the TF coils (a higher fidelity and slower calculation) gives lower ripple values for the same geometry, meaning that the faster single-filament Biot-Savart model used in the TF shape optimisation procedure is conservative.

This study was performed using a surrogate model in PROCESS of the maximum vertically stable elongation, giving κ_{95} as a function of A and m_s , see [110]. Here, a value of $m_s = 0.2$ was used, and as such we would expect the values of m_s calculated in BLUEMIRA to be approximately equal to this. However, as can be seen in figure 28, m_s is found to be higher than was accounted for in PROCESS across much of the space explored, meaning that potentially higher plasma elongation is achievable for the same passive vertical stability margin. It’s also clear that the dependency of m_s on A is not fully captured, and that it is likely that the surrogate model in PROCESS for the maximum vertically stable κ_{95} should include a dependency on R_0 as the shallow minima in m_s shown in figure 28 correlate well with the shallow minima in R_0 shown in figure 23.

There is a significant amount of noise in the values of m_s shown in figure 28, which is in large part due to the positioning of the PF coils and their currents, i.e. the destabilising term in the denominator in equation (11). This is because the coils are positioned without any stability criteria in mind and the fact that the coil position optimisation problem formulation used here is ill-determined. Note that accounting for the inner vacuum vessel shell and/or removing the passive conductors where the ports are located can drastically affect the values of m_s for a given machine configuration, and without going into detail here on the electro-mechanical attachments of port plugs we consider these results indicate that there is potential to increase κ_{95} in a large portion of the $\{A, n_{TF}\}$ space explored, thus potentially reducing the size of the machine due to improved plasma performance at higher elongation.




Figure 24. 2-D poloidal cross-section and 3-D CAD geometries of the reactor design points at $A = 2.6, 3.1$, and 3.6 , with $n_{TF} = 16$.

6.1.2. Run-time and failure rates. The total run-time for the points scanned using the full procedure outlined in section 5 averages 8.7 mins¹³. The distribution of the PROCESS and BLUEMIRA run-times is shown in figure 29.

The failure rate of the PROCESS runs was 1.25%, and the failure rate of the BLUEMIRA runs was 8.25%. The failures in the PROCESS runs were due to the gradient-based optimiser

not being able to converge to a solution satisfying the constraints from the specified starting point. The failures in the BLUEMIRA runs were most commonly in the creation and running of the CSG neutronics model, as well as in parametric geometry operations, where some boolean or offsetting operations occasionally fail due to precision issues. This is discussed further in section 8.1.

6.2. Taking a design further

There are two general paths in which the results presented the previous section can be addressed.

¹³ Timings were obtained using a 16-core machine with 1st-generation Intel® Core™ Ultra 155 H [118] processors. Only the OpenMC CSG model calculations were carried out on multiple cores.




Figure 25. Comparison of the PROCESS and BLUEMIRA results for $P_{el,net}$ for the considered range of aspect ratios at $n_{TF} = 16$.

One approach would be to address the inconsistencies between the different levels of fidelity, e.g. addressing PROCESS' inaccurate estimation of $P_{el,net}$, $\delta_{TF,peak}$, etc. This would involve modifying some design stages, possibly creating some surrogate models using the results at higher fidelity (e.g. replacing the surrogates for the peak field and ripple, or adding a function in PROCESS that calculates the fraction of power deposited in the divertor for a given A), or simply changing some inputs (e.g. changing the input energy multiplication factor). The scan could then be performed again to determine a new optimal A .

The other approach would be to proceed with the definition of design point, accepting that the initial results were imperfect and favouring instead to explore an indicative design point in more depth. This might involve fixing most inputs, performing some design steps again with higher levels of discretisation (e.g. for equilibria), performing some further loosely coupled analyses (e.g. fuel cycle assessment), improving some aspects of the automated geometrical design of the reactor manually and issuing some volumes as space reservations for more detailed design and analysis (e.g. blankets). In order to progress a conceptual design point generated automatically via a parameterisation in BLUEMIRA into the engineering design phase, manual intervention is a wise course of action. The extent of this effort, however, is far lower than would have otherwise been necessary.

For the EU-DEMO and STEP design parameterisations in BLUEMIRA, a pragmatic mixture of both approaches is followed, usually in combination with some margins on parameters known to be poorly estimated in PROCESS.

7. Case study: exploring alternative design choices

In this section, using the same reactor parameterisation described in section 5, we explore some alternative design decisions and sub-system parameterisations.

7.1. Alternative TF coil shape parameterisations

Here, we demonstrate alternative geometry parameterisations for the TF coil poloidal shape. We use the aforementioned design workflow, and solve for the same design optimisation problem (see equation (21)), at $A = 3.125$ and $n_{TF} = 16$, but use three different shape parameterisations: (i) triple-arc (see figure 9), (ii) poly-spline (see figure 9), and a bending-free shape. The bending-free shape is solved for using an iterative numerical procedure in which the tension along the coil CCL is kept constant. The three resulting machine designs are shown in figure 30.

These three shapes have different benefits and drawbacks. The triple-arc shape (the reference shape used in the previous scans) is easier to manufacture on account of there being a finite number of circular arcs to wind along, whereas the other two exhibit constantly varying curvature (except for the straight inner leg). The poly-spline shape yields the smallest TF coil shape, and thereby the lowest stored energy of the three. The bending-free shape naturally has the lower in-plane forces, but is significantly larger (higher stored energy), and also causes the PF coils to be further away, complicating the design of the PF system. In this configuration with the bending-free TF coil shape, the pulse length constraint of 2 h cannot be met.

7.2. Alternative blanket designs

We briefly explore alternative blanket concepts in this design parameterisation, comparing the reference case (HCCB blanket with 60% Li-6 enrichment [116]), to the same material composition but with varying enrichment levels and a dual-cooled lithium-lead blanket (DCLL) [119] over a range of aspect ratios, see figure 31.

The energy multiplication does not vary significantly with aspect ratio, whereas the TBR and power deposition in the blanket (as seen by the trend in $P_{el,net}$) decrease with increasing aspect ratio, due to higher relative volumetric fraction occupied by the divertor, as discussed in section 6.1. The average values of e_{mult} and TBR compare reasonably well with the values from the literature from which the blanket poloidal discretisation and material composition were approximated (for HCPB, $e_{mult} = 1.35$ and TBR = 1.20 [116], and for DCLL, $e_{mult} = 1.21$ and TBR = 1.196 [119]).

8. Discussion and further work

8.1. Lack of robustness in parametric procedures

Parametric procedures in general, but in particular when dealing with geometry, can exhibit a certain propensity for failure.




Figure 26. Comparison of the PROCESS and BLUEMIRA results for the peak field in the TF coil winding pack for the considered range of aspect ratios and n_{TF} ranging from 15 (left) to $n_{TF} = 18$ (right).




Figure 27. Comparison of the PROCESS and BLUEMIRA results for $\delta_{TF,peak}$ for the considered range of aspect ratios and n_{TF} ranging from 15 (left) to $n_{TF} = 18$ (right). The PROCESS result is calculated retrospectively for a single Biot-Savart filament per TF coil, using the TF coil shape provided by PROCESS. Two BLUEMIRA results are shown for a ‘triple-arc’ coil shape: for a single Biot-Savart filament per TF coil at a low number of points (10) on the reference LCFS (used in the TF shape optimisation constraint function); and for a rectangular cross-section current source at a high number of points (50) on the reference LCFS.

CAD geometry operations such as offsetting wires or surfaces, sweeping, lofting, fitting splines, and Boolean operations (e.g. union, cut, fuse) can fail due to geometry precision issues. In BLUEMIRA, some fallback operations in the case of spline or offset failures are available but routinely lead to large polygons, which themselves can cause failure in subsequent geometry operations. It is suspected that a large portion of these failures arise from the use of the FreeCAD [69] wrapper around OpenCASCADE [70]. We are presently investigating the use of CadQuery [120] as an alternative wrapper, in the

hopes that issues with certain problematic operations can be resolved.

The axisymmetric CSG model is also prone to failure, and is responsible for the vast majority of failed runs in the results shown in section 6.1. This is due to way in which the blanket and divertor geometry is discretised in CSG, in which invalid CSG geometry is created, chiefly because the discretisation is designed for fast run-times and relies on the convexity of certain surfaces which can occasionally be made non-convex because of some algorithms designed to preserve volumes




Figure 28. Comparison of the BLUEMIRA results for m_s to the PROCESS input value for the considered range of aspect ratios and n_{TF} ranging from 15 (left) to $n_{TF} = 18$ (right). A second order polynomial trend line is fit through the data using the RANSAC algorithm [117] which discounts outliers.




Figure 29. Distributions of the PROCESS and BLUEMIRA (excluding PROCESS) run-times for the successful design points over the $\{A, n_{TF}\}$ space explored.

whilst preventing overlapping geometry. The discretisation of the blankets and divertors in particular to a fairly minimal set of CSG cells is also, by nature, sensitive to small changes in geometry (particularly the angles of the plasma-facing cells which ultimately determine the poloidal plane geometry in our procedure). This to some extent explains the noise in the neutronics results (alongside the relatively low number of neutrons simulated).

Finally, parametric design optimisation problems do not always succeed. They may not have valid solutions for some parts of the design space or, if a gradient-based optimiser is used (as is often the case because they are faster) a poor local minimum may be returned. The optimiser termination criteria (e.g. absolute or relative tolerances on the change in objective function or residuals), the variable bounds, and the constraints and their tolerances may also not be appropriate over the full range of the parameter space explored. Occasionally round-off errors can occur, where the optimisation is terminated due to precision issues, but may still return a useful result. In BLUEMIRA by default we assume that a round-off termination or termination by maximum number of allocated evaluations (common in gradient-free algorithms) are valid ‘good’ (but not optimal) solutions, but this is not always the case and is often not trivial to determine. This can lead to failures in subsequent procedures or outliers in certain results.

8.2. Presently untreated technical aspects

Whilst a large variety of typical tokamak design activities are to some extent supported in BLUEMIRA, there remain several aspects that are poorly covered or ignored altogether. In particular, BLUEMIRA places relatively little emphasis on the developments and the analysis of the plasma scenario. Some simplified confinement, stability, radiative, and experimentally derived density and pressure limits are covered by e.g. PROCESS, but are not assessed in further depth in BLUEMIRA. Aspects such as heating and current drive, disruption mitigation, ramp-up and ramp-down, and the vast majority of transient plasma events and, indeed, engineering safety cases (e.g. loss of coolant accidents, earthquakes, etc) are not addressed.

We believe that many of the physics aspects, including scenario development and analysis of plasma transients, are




Figure 30. Alternative TF coil poloidal shape parameterisations for a design point at $A = 3.125$, $n_{TF} = 16$.




Figure 31. Energy multiplication, TBR, and net electric power as a function of A for HCCB (with varying Li-6 enrichment levels) and DCLL blanket concepts.

best treated in dedicated plasma ‘flight simulators’, such as FENIX [121, 122]. Results from such dedicated studies could be treated in simplified ways in both traditional systems codes and BLUEMIRA by means of e.g. input values, or additional constraints.

In many cases, detailed engineering analyses of transient and/or critical safety cases (which can often be design-driving) require detailed descriptions of geometry and load cases, and are computationally expensive. The aim of the BLUEMIRA code is to reach a design point where the inputs for such studies can be delivered with relative ease.

8.3. Mid-term technical roadmap

An issue with any sequential design procedure where some results are progressively ‘over-written’ by higher fidelity estimates is that inconsistencies with the original design basis are built up as the design point definition progresses, see e.g. section 6.1.1. This is usually the case with tokamak design procedures, in BLUEMIRA or elsewhere. An outstanding issue in BLUEMIRA is to provide the infrastructure to address these ‘fidelity’ inconsistencies, by means of e.g. surrogate models or fixed-point iteration. Such infrastructure would be the

foundation of any approach to a tractable and unsupervised, full-reactor optimisation scheme, where the evaluation function would be a single pass through a self-consistent sequential design procedure (with some associated objective and constraint functions). This is a key limitation of such ‘multi-fidelity’ workflows and directly impacts on the validity of design exploration studies performed in this manner. Whilst many manual design workflows exhibit the same issues, automated workflows are better positioned to actively address them, and this will be the subject of future work.

The 1.5-D transport and fixed boundary solver, PLASMOD [41, 42], is at present the only feature in BLUEMIRA which is not open-source. In future, this may be replaced by JETTO [45], which may be open-sourced in future, or TORAX [123] which is already open-source.

Improvements to both the CSG CAD neutronics pipeline (in terms of robustness and flexibility) and the 3-D CAD neutronics pipeline (in terms of surface mesh quality) are planned.

A transient full-pulse electrical power balance model, including active and reactive loads, is being developed and will be included in future releases.

Whilst an interface to the Fenics-X [55] finite element framework does exist (not described here), it is presently only functional in 2-D. In future, the automated meshing and finite element analysis of geometries created in BLUEMIRA will be extended to 3-D, including the parametric specification of boundary conditions, and loads, and material properties.

One important aspect of further work will be to include an easily configurable cost model for a given reactor, above and beyond what is already offered in PROCESS. With an automatically generated CAD model, accurate volume calculations of the reactor components are available. Fusion reactor cost models, however, have a tendency to be polemic, with a wide variety of approaches possible (e.g. top-down, bottom-up) and rely heavily on imperfect assumptions (raw material costs, manufacturing costs, interest rate, etc). Because of the sensitive nature of fusion reactor costing, it is also very likely that specific implementations of cost models will remain private. Nevertheless, we would like to offer some basic functionality to support such assessments.

8.4. Mid-term software roadmap

Despite concerted efforts, the installation procedure remains lengthier than for smaller codes. This is due to the larger number of relatively heavy dependencies, and the need to ensure that they can all operate in the same environment. The installation of OpenMC along with the download and pre-processing of the nuclear library data is optional, but is required to perform neutronics analyses in BLUEMIRA.

The past years of development have focussed primarily on creating a stable and modular software architecture and delivering a relatively broad set of functionality in terms of physics and engineering models. This focus has, however, detracted somewhat from the perspective of the front-end user. At present, there are notable user interface and user experience challenges which in large part stem from a high degree of configurability and our best efforts to try and make models and

workflows as generic as possible. These user experience issues will be addressed in the near future, with the development of a more graphical and intuitive user interface for existing reactor parameterisations, and more support for developing custom design parameterisations.

9. Summary

A modular, open-source tokamak fusion reactor design framework, BLUEMIRA, is presented. A wide variety of modules covering typical tokamak design activities ranging from systems code runs and plasma equilibrium design to 3-D CAD generation and neutronics analyses are included. Custom reactor design parameterisations and workflows can be assembled by sequentially combining a series of design stages, either using or adapting existing design and build stages, or writing one’s own.

We demonstrate the application of the BLUEMIRA framework to a conventional aspect ratio tokamak reactor design, describing the parameterisation of the design, and reporting on an indicative design space exploration study in aspect ratio and number of TF coils. The average run-time for each of the design points, from systems code run, through coil cage design, CAD generation, and CSG neutronics is well under 10 mins on a modern laptop. BLUEMIRA was able to characterise the design space in far higher fidelity than PROCESS, orders of magnitude faster and cheaper than could have been achieved using a conventional manual workflow.

We invite those interested in using and/or contributing to this framework to get in touch via the online repository [25].

Acknowledgment

We would like to thank M. Kovari (UKAEA), Dr M. Siccino (IPP), Prof. H. Zohm (IPP), Dr E. Fable (IPP), Dr R. Kembleton (UKAEA), Prof. M. Gorley (UKAEA), Dr A. Harte (UKAEA), Dr A. J. Pearce (UKFE), Dr A. Davis (UKAEA), C. Ashe (UKFE), and T. Nunn (UKAEA) for their support over the years, and all the contributors to the BLUEMIRA repository [25], and the other tools presented here [31–33], and [35].

This work has been part-funded by the EPSRC Fusion Grant 2022/27 [grant number EP/W006839/1]. To obtain further information on the data and models underlying this paper please contact PublicationsManager@ukaea.uk.

For the purpose of open access, the author(s) has applied a Creative Commons Attribution (CC BY) licence (where permitted by UKRI, ‘Open Government Licence’ or ‘Creative Commons Attribution No-derivatives (CC BY-ND) licence’ may be stated instead) to any Author Accepted Manuscript version arising.

This work has been part-funded by the Fusion Futures Programme. As announced by the UK Government in October 2023, Fusion Futures aims to provide holistic support for the development of the fusion sector.

This work has been carried out within the framework of the EUROfusion Consortium, funded by the European Union

via the Euratom Research and Training Programme (Grant Agreement No 101052200 — EUROfusion) and from the EPSRC [grant number EP/W006839/1]. To obtain further information on the data and models underlying this paper please contact PublicationsManager@ukaea.uk. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the European Commission can be held responsible for them.

ORCID iDs

M. Coleman  0000-0003-0557-6036

F. Franza  0000-0002-6950-0447

I.A. Maione  0000-0003-0224-9195

H. Baker  0009-0002-8457-7490

A. Nilima  0000-0001-9606-2676

D. Vaccaro  0009-0007-2171-794X

References

- [1] Franza F. 2019 Development and validation of a computational tool for fusion reactors system analysis (<https://doi.org/10.5445/IR/1000095873>)
- [2] Franza F., Boccaccini L., Fable E., Landman I., Maione I., Petschanyi S., Stieglitz R. and Zohm H. 2022 MIRA: a multi-physics approach to designing a fusion power plant *Nucl. Fusion* **62** 076042
- [3] Coleman M. and McIntosh S. 2019 BLUEPRINT: a novel approach to fusion reactor design *Fusion Eng. Des.* **139** 26–38
- [4] Coleman M. 2021 An integrated design framework for future nuclear fusion power reactors *Thesis* University of Cambridge, 2021-10-27T02:20:15Z
- [5] Meneghini O. *et al* 2024 FUSE (fusion synthesis engine): a next generation framework for integrated design of fusion pilot plants (arXiv:2409.05894 [physics])
- [6] Slendebroek T. *et al* 2026 Exploring the fusion power plant design space: comparative analysis of positive and negative triangularity tokamaks through optimization *Nucl. Fusion* **66** 026032
- [7] Collins C., Park J.M., Barnett R., Borowiec K., Hassan E., Humrickhouse P., Kim K., Lore J., Snyder P. and Badalassi V. The Freda Team 2023 Introducing FREDa: an integrated physics and engineering tool for fusion reactor design, assessment and optimization *APS Division of Plasma Physics Meeting Abstracts, Vol. 2023 of APS Meeting Abstracts* p CO09.005
- [8] Federici G. *et al* 2019 Overview of the DEMO staged design approach in Europe *Nucl. Fusion* **59** 066013
- [9] Wilson H., Chapman I., Denton T., Morris W., Patel B., Voss G. and Waldon C. (STEP Team) 2020 *Step—on the Pathway to Fusion Commercialization, in: Commercialising Fusion Energy: How Small Businesses are Transforming big Science* (IOP Publishing)
- [10] Song Y. *et al* 2022 Engineering design of the CFETR machine *Fusion Eng. Des.* **183** 113247
- [11] Tobita K., Utoh H., Hiwatari R., Miyoshi Y., Tokunaga S., Sakamoto Y., Someya Y., Asakura N., Homma Y. and Nakajima N. 2019 Conceptual design of Japan's fusion DEMO reactor (JADEMO) and superconducting coil issues *J. Phys.* **1293** 012078
- [12] Kim K. *et al* 2015 Design concept of K-DEMO for near-term implementation *Nucl. Fusion* **55** 053027
- [13] Menard J. *et al* 2016 Fusion nuclear science facilities and pilot plants based on the spherical tokamak *Nucl. Fusion* **56** 106023
- [14] Buttery R., Park J., McClenaghan J., Weisberg D., Canik J., Ferron J., Garofalo A., Holcomb C., Leuer J. and Snyder P. 2021 The advanced tokamak path to a compact net electric fusion pilot plant *Nucl. Fusion* **61** 046028
- [15] El-Guebaly L.A. 2010 Fifty years of magnetic fusion research (1958–2008): Brief historical overview and discussion of future trends *Energies* **3** 1067–86
- [16] Sorbom B.N. *et al* 2015 ARC: A compact, high-field, fusion nuclear science facility and demonstration power plant with demountable magnets *Fusion Eng. Des.* **100** 378–405
- [17] Rodriguez-Fernandez P. *et al* 2022 Overview of the SPARC physics basis towards the exploration of burning-plasma regimes in high-field, compact tokamaks *Nucl. Fusion* **62** 042003
- [18] White S. 2024 First design details of fusion energy pilot plant revealed by Tokamak Energy at top physics conference
- [19] Rickover A.H.G. 1953 letter to an unknown recipient, later recanted in part during one of Rickover's testimonies to the US Congress Paper reactors, real reactors
- [20] Quartararo A. Systems codes - DEMO LAR Gate Review 2a 2025
- [21] Davis A., Waldon C., Muldrew S.I., Patel B.S., Verrier P., Barrett T.R. and Politis G.A. 2024 Digital: accelerating the pathway *Philos. Trans. R. Soc. A* **382** 20230411
- [22] Muldrew S.I. *et al* 2024 Conceptual design workflow for the STEP prototype powerplant *Fusion Eng. Des.* **201** 114238
- [23] Vaccaro D. *et al* 2024 Models implemented in the methodological approach to design the initial STEP first wall contour *Fusion Eng. Des.* **204** 114479
- [24] Nilima A., Bluteau M., Cook J., Funk O., Graham G., Matthews J., Muldrew S.I., Pearce A.J. and Vaccaro D. 2025 Optimisation of STEP poloidal field coils with superconducting coil constraints in STEP-Bluemira power plant design framework *Fusion Eng. Des.* **220** 115357
- [25] Coleman M. *et al* 2023 bluemira
- [26] Free Software Foundation 1999 GNU lesser general public license v2.1 -GNU project
- [27] Romano P.K., Horelik N.E., Herman B.R., Nelson A.G., Forget B. and Smith K. 2015 OpenMC: a state-of-the-art Monte Carlo code for research and development *Ann. Nucl. Energy* **82** 90–97
- [28] Merkel D. 2014 Docker: lightweight linux containers for consistent development and deployment *Linux J.* **239** 2
- [29] Lao L. 1997 G-EQDSK format specification
- [30] Sauter O. and Medvedev S. 2013 Tokamak coordinate conventions *Comput. Phys. Commun.* **184** 293–302
- [31] Coleman M., Cook J.E., Franza F., Maione I.A., McIntosh S., Morris J., Short D., Blair A.I., Funk O. and Saunders H. 2025 eqdsk
- [32] Coleman M. and Cook J.E. 2025 matproplib
- [33] Coleman M., Cook J.E. and Wong O. 2017 tokamak-neutron-source
- [34] Tautges T.J., Ernst C., Stimpson C., Meyers R.J. and Merkley K. 2004 MOAB : a mesh-oriented database *Tech. Rep. SAND2004-1592* Sandia National Laboratories (SNL)
- [35] Cook J.E., Funk O. and Saunders H. 2025 fast_ctd
- [36] Bray T. 2017 The JavaScript Object Notation (JSON) data interchange format, RFC 8259
- [37] Imbeaux F. *et al* 2015 Design and first applications of the ITER integrated modelling & analysis suite *Nucl. Fusion* **55** 123006
- [38] STEP-file, ISO 10303-21 2023
- [39] Kovari M., Kemp R., Lux H., Knight P., Morris J. and Ward D.J. 2014 PROCESS: a systems code for fusion power plants—Part 1: Physics *Fusion Eng. Des.* **89** 3054–69

- [40] Kovari M., Fox F., Harrington C., Kembleton R., Knight P., Lux H. and Morris J. 2016 PROCESS[™]: a systems code for fusion power plants – Part 2: Engineering *Fusion Eng. Des.* **104** 9–20
- [41] Fable E., Angioni C., Ivanov A., Lackner K., Maj O., Yu S., Medvedev G.P. and Pereverzev G. 2013 A stable scheme for computation of coupled transport and equilibrium equations in tokamaks *Nucl. Fusion* **53** 033002
- [42] Fable E., Angioni C., Siccino M. and Zohm H. 2018 Plasma physics for fusion reactor system codes: Framework and model code *Fusion Eng. Des.* **130** 131–6
- [43] Goedbloed J. and Zakharov L. 1982 A comparison of two approaches for solving equilibrium problems for axisymmetric current-carrying plasmas, *Tech. Rep.* rijnhuizen–82-143 INIS Reference Number: 14763133
- [44] Zakharov L. 1985 Electrodynamical moments method for calculating toroidal plasma equilibrium IAE-4114/6
- [45] Cenacchi G., Taroni A. and Enea R.I. 1988 JETTO a free boundary plasma transport code *Tech. Rep.* ENEA-RT-TIB-88-5, ENEA
- [46] Zhang D. and Koh C.S. 2012 An efficient semianalytic computation method of magnetic field for a circular coil with rectangular cross section *IEEE Trans. Magn.* **48** 62–68
- [47] Babic S.I. and Akyel C. 2005 An improvement in the calculation of the magnetic field for an arbitrary geometry coil with rectangular cross section *Int. J. Numer. Modelling. Electron. Netw. Devices Fields* **18** 493–504
- [48] Babic S.I. and Akyel C. 2008 Erratum: An improvement in the calculation of the magnetic field for an arbitrary geometry coil with rectangular cross section *Int. J. Numer. Model.* **18** 493–504
- [49] Feng Z. 1985 The treatment of singularities in calculation of magnetic field by using integral method *IEEE Trans. Magn.* **21** 2207–10
- [50] Sackett S.J. 1977 EFFI: a code for calculating the electromagnetic field, force, and inductance in coil systems of arbitrary geometry. User's manual, *Tech. Rep.* UCID-17621 California Univ Lawrence Livermore National Laboratory
- [51] Grad H. and Rubin H. 1958 Hydromagnetic equilibria and force-free fields, *Tech. Rep.* A/CONF.15/P/386, New York Univ. Inst. of Mathematical Sciences
- [52] Shafranov V.D. 1966 Plasma Equilibrium in a Magnetic Field *Rev. Plasma Phys.* **2** 103
- [53] Geuzaine C., Remacle J.-F. and G. 2009 Gmsh: A 3-D finite element mesh generator with built-in pre- and post-processing facilities *Int. J. Numer. Methods Eng.* **79** 1309–31
- [54] Logg A., Mardal K.-A. and Wells G. 2012 Automated Solution of Differential Equations by the Finite Element Method *The Fenics Book*, vol 84 of (*Lecture Notes in Computational Science and Engineering*) (Springer)
- [55] Baratta I.A., Dean J.P., Dokken J.S., Habera M., Hale J.S., Richardson C.N., Rognes M.E., Scroggs M.W., Sime N. and Wells G.N. 2023 DOLFINx: the next generation FEniCS problem solving environment
- [56] Zheng S.B., Wootton A.J. and Solano E.R. 1996 Analytical tokamak equilibrium for shaped plasmas *Phys. Plasmas* **3** 1176–8
- [57] Coleman M. and McIntosh S. 2020 The design and optimisation of tokamak poloidal field systems in the BLUEPRINT framework *Fusion Eng. Des.* **154** 111544
- [58] Graham G., Bardsley O., Kahn S., Pearce A. and Johnson M. 2023 Conceptual design of spherical tokamaks using spherical harmonic flux constraints in Bluemira *Proc. 30th IEEE Symp. on Fusion Engineering*
- [59] Bardsley O.P., Baker J.L. and Vincent C. 2024 Decoupled magnetic control of spherical tokamak divertors via vacuum harmonic constraints *Plasma Phys. Control. Fusion* **66** 055006
- [60] Mould C. et al 2025 Implementing toroidal harmonic flux constraints in Bluemira *Open-Source Software for Fusion Energy* (available at: <https://zenodo.org/records/15053682>)
- [61] Graham G. et al 2026 Applications of spherical and toroidal harmonic equilibrium constraints in BLUEMIRA (in preparation)
- [62] Amorisco N.C., Agnello A., Holt G., Mars M., Buchanan J. and Pamela S. 2024 FreeGSNKE: a Python-based dynamic free-boundary toroidal plasma equilibrium solver *Phys. Plasmas* **31** 042517
- [63] Kraft D. 1988 A Software Package for Sequential Quadratic Programming (Wiss. Berichtswesen d. DFVLR)
- [64] Powell M.J.D. 1994 A Direct Search Optimization Method That Models the Objective and Constraint Functions by Linear Interpolation *Advances in Optimization and Numerical Analysis*, Ed S. Gomez and J.-P. Hennart (Springer) pp 51–67
- [65] Coutlis A., Bandyopadhyay I., Lister J.B., Vyas P., Albanese R., Limebeer D.J.N., Villone F. and Wainwright J.P. 1999 Measurement of the open loop plasma equilibrium response in TCV *Nucl. Fusion* **39** 663
- [66] Lister J.B., Sharma A., Limebeer D.J.N., Nakamura Y., Wainwright J.P. and Yoshino R. 2002 Plasma equilibrium response modelling and validation on JT-60U *Nucl. Fusion* **42** 708
- [67] Leuer J.A. 1989 Passive vertical stability in the next generation tokamaks *Fusion Technol.* **15** 489–94
- [68] Cunningham G. 2017 private communication
- [69] (FreeCAD) 2019 FreeCAD source documentation (available at: www.freecadweb.org)
- [70] OPEN CASCADE S.A.S., Open CASCADE 2015
- [71] Shimwell J. et al 2019 Multiphysics analysis with cad-based parametric breeding blanket creation for rapid design iteration *Nucl. Fusion* **59** 046019
- [72] Gralnick S.L. and Tenney F.H. 1975 Analytical solution of the toroidal constant tension solenoid, *Tech. Rep.* CONF-751125-122, Princeton Univ., N.J. (USA). Plasma Physics Laboratory
- [73] Manickam J. 1984 Stability of $n = 1$ internal modes in tokamaks *Nucl. Fusion* **24** 595
- [74] Johner J. 2011 HELIOS: a Zero-Dimensional Tool for Next Step and Reactor Studies *Fusion Sci. Technol.* **59** 308–349
- [75] Kuiroukidis A. and Throumoulopoulos G.N. 2015 An analytic functional form for characterization and generation of axisymmetric plasma boundaries (2013 Plasma Phys. Control. Fusion 55 095009) *Plasma Phys. Control. Fusion* **57** 078001
- [76] Rising M.E. et al 2025 MCNP[®] Code Version 6.3.1 Release Notes, *Tech. Rep.* LA-UR-25-23548, Rev. 1, Los Alamos National Laboratory
- [77] Koning A., Rochman D., Sublet J.-C., Dzysiuik N., Fleming M. and van der Marck S. 2019 Tendl: Complete nuclear data library for innovative nuclear science and technology *Nucl. Data Sheets* **155** 1–55 special Issue on Nuclear Reaction Data
- [78] Chadwick M. et al 2011 ENDF/B-VII.1 nuclear data for science and technology: Cross sections, covariances, fission product yields and decay data *Nucl. Data Sheets* **112** 2887–996 special Issue on ENDF/B-VII.1 Library
- [79] Shimwell J. 2025 openmc_data, original-date: 2024-01-27T21:23:38Z
- [80] Giannini L., Boso D.P. and Corato V. 2022 A combined electromagnetic and mechanical approach for eu-demo toroidal field coils *Appl. Sci.* **12** 2766
- [81] Fausser C., Puma A.L., Gabriel F. and Villari R. 2012 Tokamak D-T neutron source models for different plasma

- physics confinement modes *Fusion Eng. Des.* **87** 787–92
- [82] Delaporte-Mathurin R., Shimwell J., Pattinson L., Pranto A.I. and Faisal M. 2023 OpenMC Plasma Source, T09:: 57Z
- [83] Nobre G. et al ENDF/B-VIII.1
- [84] Bosch H.-S. and Hale G.M. 1992 Improved formulas for fusion cross-sections and thermal reactivities *Nucl. Fusion* **32** 611
- [85] Bosch H.-S. and Hale G.M. 1993 Improved formulas for fusion cross-sections and thermal reactivities *Nucl. Fusion* **33** 1919
- [86] Ballabio L., Källne J. and Gorini G. 1998 Relativistic calculation of fusion product spectra for thermonuclear plasmas *Nucl. Fusion* **38** 1723–35
- [87] Appelbe B. and Chittenden J. 2016 The effects of ion temperature on the energy spectra of $T + T \rightarrow 2n + \alpha$ reaction products *High Energy Density Phys.* **19** 29–37
- [88] Mouginot B. et al 2021 DAGMC - direct accelerated geometry monte carlo toolkit
- [89] The CGAL Project 2024 *CGAL User and Reference Manual* 6.0.1 edn (CGAL Editorial Board)
- [90] Blacker T.D. et al 2016 CUBIT Geometry and Mesh Generation Toolkit 15.2 User Documentation, *Tech. Rep* SAND2016-6850R, Sandia National Laboratory (SNL-NM)
- [91] Mason S. and Nonweiler J. 2025 overlap_checker, T13:17:15Z
- [92] Davis A. 2025 occ_faceter
- [93] Intel Corporation 2021 Intel® core™ i7-11800h processor (24m cache, up to 4.60 ghz) accessed: 2025-12-15
- [94] Vaccaro D. et al 2024 Models implemented in the methodological approach to design the initial step first wall contour, to be published in Fusion Engineering and Design
- [95] Pitcher C.S. and Stangeby P.C. 1997 Experimental divertor physics *Plasma Phys. Control. Fusion* **39** 779–930
- [96] Stangeby P.C. 2000 *The Plasma Boundary of Magnetic Fusion Devices* (CRC Press)
- [97] Arter W., Riccardo V. and Fishpool G. 2014 Power deposition on tokamak plasma-facing components *IEEE Trans. Plasma Sci.* **42** 1932–42
- [98] Arter W., Surrey E. and King D.B. 2015 The SMARDDA approach to ray tracing and particle tracking *IEEE Trans. Plasma Sci.* **43** 3323–31
- [99] Giroud D.C., Meakins D.A., Carr D.M., Baciero D.A. and Bertrand M.C. 2018 CHERAB spectroscopy modelling framework
- [100] Carr M. et al 2017 *Towards Integrated Data Analysis of Divertor Diagnostics With ray-Tracing*
- [101] Summers H. and O'Mullane M. 2011 Atomic Data and Modelling for Fusion: the ADAS Project *AIP Conf. Proc.* **1344** 179–87
- [102] Coleman M., Hörstensmeyer Y. and Cismondi F. 2019 DEMO tritium fuel cycle: performance, parameter explorations and design space constraints *Fusion Eng. Des.* **141** 79–90
- [103] Day C. and Giegerich T. 2013 The direct internal recycling concept to simplify the fuel cycle of a fusion power plant *Fusion Eng. Des.* **88** 616–20
- [104] Arredondo R., Schmid K., Subba F. and Spagnuolo G. 2021 Preliminary estimates of tritium permeation and retention in the first wall of DEMO due to ion bombardment *Nucl. Mater. Energy* **28** 101039
- [105] Meschini S., Delaporte-Mathurin R., Tynan G.R. and Ferry S.E. 2025 Impact of trapping on tritium self-sufficiency and tritium inventories in fusion power plant fuel cycles *Nucl. Fusion* **65** 036010
- [106] Wolpert D. and Macready W. 1997 No free lunch theorems for optimization *IEEE Trans. Evol. Comput.* **1** 67–82
- [107] Johnson S.G. 2018 The NLOpt nonlinear-optimization package (available at: <http://ab-initio.mit.edu/nlopt>)
- [108] Virtanen P. et al 2020 SciPy 1.0: fundamental algorithms for scientific computing in Python *Nat. Methods* **17** 261–72
- [109] Federici G. et al 2016 Overview of the design approach and prioritization of R&D activities towards an EU DEMO *Fusion Eng. Des.* **109–111** 1464–74
- [110] Coleman M., Zohm H., Bourdelle C., Maviglia F., Pearce A., Siccino M., Spagnuolo A. and Wiesen S. 2025 Definition of an EU-DEMO design point robust to epistemic plasma physics uncertainties *Nucl. Fusion* **65** 036039
- [111] Loving A. et al 2014 Pre-conceptual design assessment of DEMO remote maintenance *Fusion Eng. Des.* **89** 2246–50
- [112] Crofts O., Loving A., Iglesias D., Coleman M., Siuko M., Mittwollen M., Queral V., Vale A. and Villedieu E. 2016 Overview of progress on the European DEMO remote maintenance strategy *Fusion Eng. Des.* **109–111** 1392–8
- [113] Albanese R., Ambrosino R., Castaldo A. and Loschiavo V. 2018 Optimization of the PF coil system in axisymmetric fusion devices *Fusion Eng. Des.* **133** 163–72
- [114] Albanese R., Ambrosino R. and Mattei M. 2015 CREATE-NL+: a robust control-oriented free boundary dynamic plasma equilibrium solver *Fusion Eng. Des.* **96–97** 664–7
- [115] Pereslavytsev P., Hernández F.A., Zhou G., Lu L., Wegmann C. and Fischer U. 2019 Nuclear analyses of solid breeder blanket options for DEMO: Status, challenges and outlook *Fusion Eng. Des.* **146** 563–7
- [116] Zhou G., Hernández F.A., Pereslavytsev P., Kiss B., Rethesh A., Maqueda L. and Park J.H. 2023 The European DEMO Helium Cooled pebble bed breeding blanket: design status at the conclusion of the pre-concept design phase *Energies* **16** 5377
- [117] Fischler M.A. and Bolles R.C. 1981 Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography *Commun. ACM* **24** 381–95
- [118] Intel Corporation 2023 Intel® Core™ Ultra 7 Processor 155H (24M Cache, up to 4.80 GHz) - Product Specifications
- [119] Palermo I., Fernández-Berceruelo I., Rapisarda D. and Ibarra A. 2019 Neutronic assessments towards a comprehensive design of DEMO with DCLL breeding blanket *Fusion Eng. Des.* **138** 217–25
- [120] (Cadquery Contributors) (CadQuery)
- [121] Fable E. et al (the ASDEX Upgrade Team) 2022 The modeling of a tokamak plasma discharge, from first principles to a flight simulator *Plasma Phys. Control. Fusion* **64** 044002
- [122] David P., Fable E., Wu C., Gehring M., Sieben M., Sieglin B., Tardini G. and Treutterer W. 2018 Update of the integrated flight simulator for ASDEX Upgrade, Open Plasma Science
- [123] Citrin J. et al 2024 TORAX: a fast and differentiable tokamak transport simulator in JAX (arXiv:2406.06718 [physics])