

BLUEPRINT: A novel approach to fusion reactor design

M. Coleman^{a,b,*}, S. McIntosh^{a,c}

^a United Kingdom Atomic Energy Authority, Culham Science Centre, Abingdon, Oxfordshire OX14 3DB, United Kingdom

^b Department of Engineering, University of Cambridge, Trumpington Street, Cambridge CB2 1PZ, United Kingdom

^c ITER Organization, Route de Vinon-sur-Verdon, CS 90 046, 13067 St Paul Lez Durance Cedex, France



ARTICLE INFO

Keywords:

DEMO
Fusion reactor
Systems code
Design

ABSTRACT

The European DEMO fusion power reactor (EU-DEMO) is still in the pre-conceptual design phase. The design strategy for the EU-DEMO hinges on investigating multiple reactor designs and technologies in parallel, progressively down-selecting these in the mid-2020s, in preparation for the conceptual design phase. The present implementation of the strategy centres around a baseline single-null design, which is configuration controlled and iterated approximately every two years. The majority of resources are dedicated to studying the baseline design; changing it is expensive, and takes months to do. Competing technologies for different sub-systems (e.g. blankets, magnets) are forced to co-exist within the same physical confines (CAD models), and conform to the same set of performance criteria. Meanwhile, the alternative reactor designs are only loosely defined, with no agreed set of parameters, no CAD models, and no formal framework for study.

We argue that the EU-DEMO design strategy is best served by a more efficient implementation methodology which would enable more comprehensive exploration of the DEMO design space. We make the case for a change in approach to the design of DEMO-class reactors in Europe, and propose a solution which bears some resemblance to the present methodology, yet automates and accelerates the sequence of design and analysis activities when generating a design configuration.

We present our preliminary attempts to demonstrate the feasibility of our idea, in the form of a new fusion reactor design code: BLUEPRINT. We demonstrate that the typical activities required to generate a DEMO design point can be sped up by *four orders of magnitude*—from months to minutes—paving the way for a rigorous and broad exploration of the design space.

1. Introduction

Throughout the last decade, European efforts in fusion power reactor R&D have been coordinated under the European Fusion Development Agreement (EFDA) and its successor, EUROfusion, following the programme and strategy laid out in the European Roadmap to Fusion Electricity [1]. The power plant design and R&D activities carried out in the European Union (EU) under this framework are vital to steer the broader fusion R&D programme and maintain a pragmatic approach to the development and realisation of fusion power.

The strategy for the design of the EU-DEMO is to investigate multiple different reactor concepts (single-null, double-null, flexi-DEMO, etc.) and technology options (blanket designs, toroidal field coil designs, etc.) and progressively down-select these, before converging upon a single reactor design with a reduced set of sub-variants in the mid-2020's, preparing for a conceptual design review in 2027 [2] (see Fig. 1).

A broad collaborative environment now exists between fusion laboratories across Europe working together on reactor design, with work in the field of power plant physics and technology focussed on the EU-DEMO reactor designs. Work is centrally coordinated by the EUROfusion Power Plant Physics and Technology (PPPT) Department, which distributes resources to various specialist work packages, each managed by a leading expert from one of the national fusion laboratories within Europe [2–4]. This European-wide effort is the largest and most comprehensive of its kind in the world today, although Chinese work on the Chinese Fusion Engineering Test Reactor (CFETR) [5] may soon challenge this.

This paper is organised in two main parts. In the first (Section 2), we describe the present strategy and approach to the design of DEMO-class reactors in the EU, and discuss its shortcomings. In the second half (Section 3), we suggest an alternative approach and carry out a proof-of-principle study to demonstrate its feasibility. The result of this proof-of-principle study is a novel reactor design code: BLUEPRINT.

* Corresponding author at: United Kingdom Atomic Energy Authority, Culham Science Centre, Abingdon, Oxfordshire OX14 3DB, United Kingdom.

E-mail address: matti.coleman@ukaea.uk (M. Coleman).

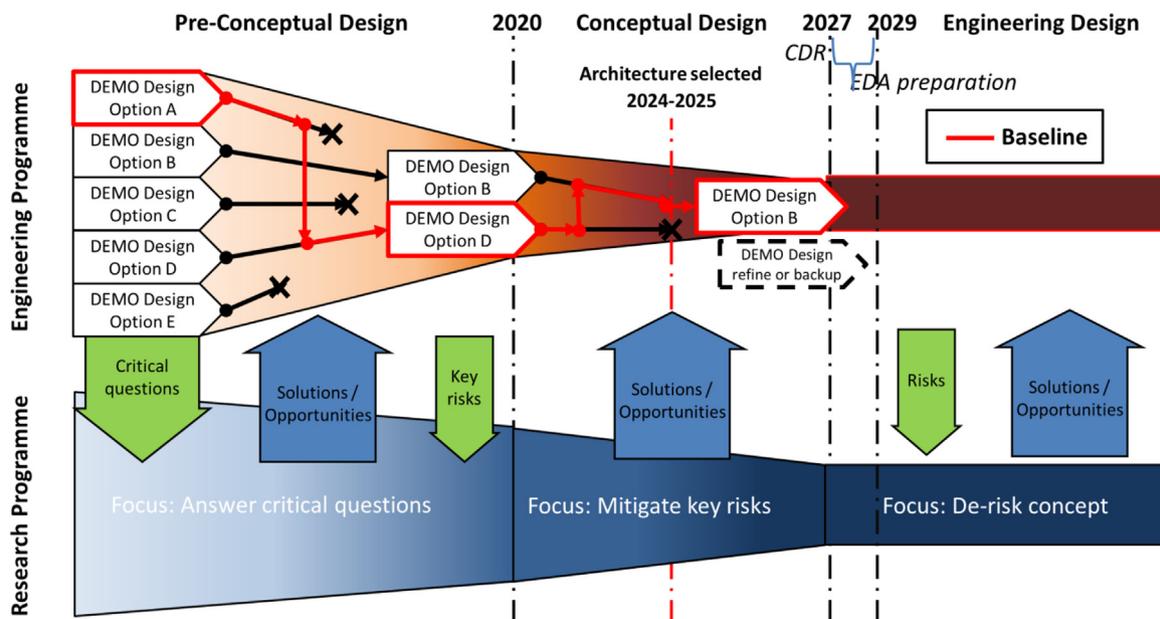


Fig. 1. EU-DEMO design strategy leading up to the start of an engineering design phase in the late 2020s [2].

This work is not a DEMO design study; it is rather a study of the DEMO design process and how it may be improved upon. In presenting our work, we focus on a single null reactor for illustrative purposes—it also happens to be the easiest to describe (due to experience and data). We do not present a novel or better reactor design, and we do not explicitly present any results pertaining to the reactor designs illustrated herein. Where necessary we assign values to relevant variables or show indicative results. This is done exclusively with the intention of facilitating understanding and/or demonstrating capability. We wish to remain agnostic to reactor designs and their respective performances. We present our methodology supported by proof-of-principle experiments only to demonstrate its feasibility. The key result is the approximate runtime of our reactor design code.

2. Present fusion reactor design approach

2.1. Strategy

The present EU strategy for the design of a DEMO-class reactor is to develop multiple alternative solutions for the same set of high-level plant requirements (such as net electric output for minimal cost), see Fig. 1. One of these alternatives, termed the “baseline”, is investigated in more detail from both a physics and engineering perspective; the rationale being that many disciplines require more detail in order to carry out meaningful studies and provide higher level information to the programme.

The baseline is an extremely useful tool, which in the EU-DEMO has been instrumental in bringing complex design and integrational issues to light. Considering the full reactor plant and integration issues from the outset has proved important in understanding the design drivers and constraints on the various systems. Naturally though, the cost of more detailed investigations is higher, and the turnaround time of such studies is much longer. Given this, the baseline design is typically frozen for a period of two to three years to enable a stable reactor configuration to be investigated.

In parallel to the baseline activities, to which 90–95% of resources for DEMO R&D studies are typically dedicated, other reactor configurations are investigated in comparatively little detail. Notionally, these include a double-null (DN) DEMO, a snowflake (SF) divertor configuration, and a “flexi-DEMO”, the latter of which aims to achieve steady-state operation in the later phases of the reactor life [2]. In most cases,

there is no agreed set of cardinal reactor parameters or 2D geometry, let alone CAD data. Indeed, in the vast majority of cases, discipline-specific work packages responsible for the design of a reactor sub-system may not even be aware of the alternative reactor configurations or the implications that they may have on their sub-systems, and very rarely consider them in their work. For example, balance of plant studies in Europe do not consider DN or SF configurations, despite the fact that both will likely mean less high-temperature heat transported to the secondary cycle. Similarly, maintenance studies of more exotic divertor configurations such as a SF are not carried out, despite the fact that they may well be incompatible with the required positions of divertor poloidal field coils. Notable exceptions include heating and current drive (H&CD) studies for steady-state reactor operation (e.g. [6,7]), and remote maintenance assessments of a DN reactor (e.g. [8]).

Recently, the EU-DEMO design programme has adopted a more focussed approach (see [9]), decoupling the baseline design point definition from so-called “key design issues”, in which alternative solutions are investigated, addressing a particular design problem (e.g. water vs. helium as a primary coolant). This reduces the need for maintaining parallel reactor designs to some extent, yet will still require detailed and wide-ranging analyses of the various alternatives being investigated for each key design issue.

2.2. Implementation

The procedure to create a baseline configuration for a DEMO reactor is a protracted one. First, a systems code is used to generate a 0-D self-consistent solution to a posed optimisation problem. This usually takes the form of a set of physics and technological constraints set in the systems code, some equality constraints (typically net electricity output and pulse length), and an optimisation objective: the minimisation of the major radius of the tokamak, as an ersatz objective for the global plant cost. Two 0/1-D systems codes capable of performing this role exist: PROCESS [10,11] and SYCOMORE [12]. Although they differ in many regards (calculation procedures, optimisers, etc.), they can be used interchangeably to generate a self-consistent reactor design point. PROCESS is normally used in EU-DEMO studies.

The actual run-times for PROCESS results are exceedingly fast; typically less than ten seconds on a single, modern computer core. However, human iterations and verification of the output are usually required, such that results can be expected within a day, once a

coherent set of inputs has been agreed.

Once these results are obtained, they are manually processed to generate a 2-D cross-section of the reactor based on the radial build output (the toroidal field (TF) coils, vacuum vessel (VV), breeding blankets (BB), etc.), which takes two to three days.

The 0-D systems code parameter set and the 2-D cross-section are then used as an input to plasma magnetic equilibrium solvers, which optimise the positions and currents of the poloidal field (PF) coils to achieve a target plasma separatrix shape. Furthermore, preliminary assessments of the vertical stability of the plasma are also conducted, to confirm that the plasma elongation and the desired equilibria are achievable with the envisaged PF coil reactive power supplies. In total, with some iteration, such studies can take as long as two weeks to complete.

Once the positions of the PF coils are known, a 3-D CAD model of the principal tokamak components can be generated. This is done manually and, for a reasonably detailed CAD representation of most of the major reactor components, can take approximately one month.

For some areas of work, the CAD model is sufficient to begin more detailed studies. However, in order to confirm that the configuration is feasible to the first order, more detailed studies are usually performed prior to a general release of a baseline configuration. The critical path activity at this stage is the creation of the global tokamak neutronics model, which is dependent on the 3-D CAD input and can take two to three months from start to finish for a full set of outputs (confirmation of the tritium breeding ratio (TBR) and shielding performance, heat deposition, etc.). Once these studies are completed, the baseline reactor data are released to the broader community across Europe, who commence design and R&D activities.

This process is summarised in Fig. 2, where we portray it as a design cycle: the design point is defined, distributed to dedicated teams, who analyse it, study specific aspects, and sporadically feed back information to the earlier stages of the design process.

The sequential activities leading up to the definition of a design point can take between four and six months to complete, due to the durations of the tasks mentioned above, the inevitable bugs in codes, oscillations in opinions, and resource availability issues.

The materials R&D aside, the work carried out in these various PPPT work packages—be it the design of superconducting magnets or breeding blankets, remote maintenance studies, or related technology R&D—is (by its very nature) far more detailed and is often tied to hardware development cycles. Consequently, such studies can take years to complete, and only sporadically feed information back into the 0-D systems codes. This information can in turn affect the whole design of the reactor, e.g. the maximum neutral beam injection (NBI) energy or the yield strength of a cryogenic steel.

2.3. Implementation issues

Regardless of the root causes, the current system presents a number of shortcomings.

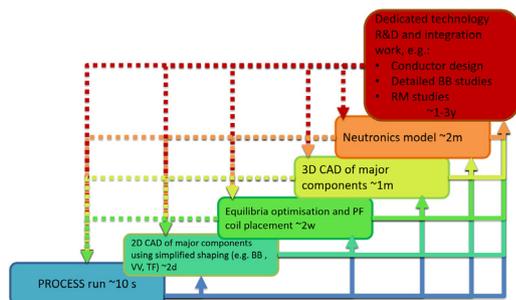


Fig. 2. Schematic representation of the DEMO reactor design point definition and evaluation procedure in Europe.

2.3.1. Cost, schedule, and coordination

Firstly, the cost and duration of these design point definition activities is very high. Due to the organisational nature of EUROfusion, and the geographical distribution of many of the underlying expertise and codes, the aforementioned calculation procedures are carried out by experts in different institutions across Europe, and coordinated centrally. The coordination and streamlining of the different activities require careful synchronisation of individual satellite actors. These actors perform tasks that are all sequentially on a critical path, which may be reset back to the starting point (systems code run) once or twice in any given cycle. It follows that the desired output cannot be expected in a timely manner, and there is a high coordination overhead associated with organising these activities.

2.3.2. Inconsistencies

Secondly, the process is prone to generating inconsistencies. If a technical or design issue is discovered late on in the definition process (e.g. during the creation of the CAD or neutronics models), there is already a considerable cost sunk in creating a certain reactor design point. Similarly, were one to restart the process, there is a significant schedule slippage to contend with—correcting an issue with the systems code run, for instance. It is therefore perhaps understandable that the (pragmatic) response to such a problem is to simply proceed with the completion of the design point, accepting a slew of inconsistencies and managing them in a variety of ways. Typically, this might involve “decoupling” one or more sub-systems or aspects from the rest of the reactor design to some extent, effectively giving a subset of the distributed team an alternative universe to operate in, so as not to incur delays in the wider programme.

2.3.3. Difficulty to develop, learn, and innovate

Thirdly, as this design point definition procedure is only carried out once every couple of years, the learning process is impaired. In each new baseline configuration, multiple different input parameters will change—as will some of the codes, models or analysis procedures, not to mention the people involved. Consequently, it becomes difficult for any one person to diagnose the underlying cause of an issue when comparing two iterations of even an almost-identical design. For example: did the peak field in the TF coils go down in the latest design because of the changes to the divertor challenge criterion, or because of the changes to the stress limits on the cryogenic steel? Is the vertical stability of the plasma better in this configuration, or are we calculating it in a different way?

Like-for-like comparisons between design points are inherently impossible, because they exist years apart, with many changes being made simultaneously. Ideas cannot readily be explored, as feedback can only be obtained months or years later. The learning process becomes a slow and frustrating exercise in disentangling various convoluted parameters and modifications to several different designs/codes/procedures, which only produce design points every couple of years. It does not lend itself to identifying trends. It does not facilitate decision-making.

The difficulties of learning from a methodology which generates only sparse data hinder the design development process. This can manifest itself in a difficulty to innovate effectively; design ideas which are too different from past design points are too unknown and therefore too expensive to explore adequately. Future design points are inherently drawn towards past iterations where data abound and confidence is higher. This results in only small or modest deviations to the baseline being accepted sporadically, with more significant changes either postponed due to lack of information, or rejected without proper study due to lack of resources. Ultimately, the more time that is spent investigating a certain type of DEMO design, the harder it is to explore alternatives, as the amount of information required to justify a change must be similar to the amount of information which already exists for the baseline.

2.3.4. Inability to compare sub-systems adequately

Finally, design options cannot be properly compared with this approach. Even within the relatively well-defined SN reactor concept, a plethora of sub-system design options still co-exist in parallel. The present approach only allows one baseline configuration to exist at any given time, purely for pragmatic reasons: the cost of generating subtle variants and the overheads of manually coordinating different analyses for each would simply be too great. Thus, the various sub-system design options are required to co-exist within the same confines: a helium-cooled pebble bed (HCPB) blanket and a dual coolant lithium-lead (DCLL) blanket must fit into the same space allocation, and meet the same shielding and breeding performance requirements. Three TF coil winding pack concepts are forced to occupy the same space and deliver the same performance.

In fact, all concepts (blanket, TF coil, or otherwise) can offer advantages to the reactor designer, at the expense of certain disadvantages. These differences are masked as all are required to operate with the same baseline model. Trade-offs cannot be carried out meaningfully because the pros and cons of different technology options cannot be elevated properly to the reactor level. The reactor designer cannot compare the higher breeding performance of the HCPB blanket against the superior shielding performance of a water-cooled lithium-lead (WCLL) blanket, because both are constricted to an arbitrary compromise. Similarly, the superior current density of graded layer-wound TF coil winding packs cannot be traded off against the lower level of technological risk of the ITER-like pancake-wound winding packs.

2.4. A fundamental mismatch: the case for a paradigm shift

The broader strategy of pursuing multiple reactor designs and investigating alternatives is not well-served by presently existing codes. Only one reactor design (the SN baseline) is investigated to a meaningful degree. The cost of this detailed investigation and the limited resources available rule out simply carrying out the same work on multiple reactors in parallel. This effectively means that many data are unavailable for the alternative designs. Yet such information is crucial if one is to progressively narrow down the concepts being investigated in the pre-conceptual design stage. How could we ever move to an alternative design if we do not know about its parameter X or performance Y? The present baseline becomes “the devil that we know”, and the inadequate study of the alternatives becomes the reason no major change in design can be decided upon—or indeed ruled out definitively.

Even if we were to decide today on a single-null design, ignoring all alternative reactor concepts, the present implementation approach risks not delivering a robust design. For the sake of argument, let us assume that a SN design is the only viable option, and that a single design iteration takes two to three years. At best, just half a dozen iterations of a SN DEMO could be expected in more than a decade of design development and R&D—rather few, if one considers the sheer complexity of the system and the number of different technology options being studied, which necessarily interact with one another. Consider the number of combinatorial SN reactor variants, when accounting for four blanket concepts, two heating and current drive systems, and three TF coil winding pack options.

The situation is exacerbated by the fact that the present approach can result in unnecessary repetition of costly design work for systems unaffected by certain design changes. CAD models must be re-built, coordinates re-set, interfaces re-defined, finite element and neutronics analyses re-built and re-run, etc. This is at present an expensive endeavour, which is understandably shunned by those involved in favour of greater stability in working assumptions, at the price of some inconsistency in the global context.

There is a lack of adequate modelling tools to design and analyse design point to a sufficiently high degree of detail. The present approach is too slow and expensive, and occasionally leads to the

propagation of inconsistencies and errors. If the fusion community is to be serious about the strategy of multiple concepts and progressive down-selection leading into the conceptual and engineering design phase, a paradigm shift in the implementation approach is required.

BLUEPRINT is a novel fusion reactor design code that enables DEMO design and analysis activities to shift from a task basis to one of algorithmic responsibility where a particular group of experts takes responsibility for an evaluation or design procedure with a clearly defined set of objectives, inputs and outputs. The key difference is that, due to the automation of many of the set up activities, each procedure can be run at any time for a given set of inputs and generate the outputs in the pure computational time required—in many cases significantly less than a normal task duration.

We estimate that such an approach would reduce the time required to define a design point to a typical level of definition by approximately four orders of magnitude. To demonstrate the feasibility of the proposed approach, in an admittedly preliminary manner, we now introduce BLUEPRINT and some of its present capabilities.

3. An introduction to BLUEPRINT

BLUEPRINT is a fusion reactor design code which is designed to significantly reduce the definition time of a design point, to enable detailed assessments of fusion reactor power plant performance to be carried out rapidly, facilitating the comparison of different technologies or reactor concepts.

Whilst it is capable of evaluating and ultimately optimising reactor design points for a given set of inputs and design decisions, it can never be a replacement for reactor designers. Rather, it is intended to be a tool to aid the reactor designers in their development efforts. The designers are required to generate ideas, codify them, integrate them into the BLUEPRINT framework, and use the code to compare new or existing alternatives. Optimisation algorithms are very good at providing answers to well-formulated problems, but they are not yet capable of suggesting entirely new ideas; the engineer must be in the loop.

3.1. Philosophy and organisation

BLUEPRINT is modular in nature and is intended, above all else, to be a framework in which different codes or calculation procedures can be selected and combined to produce coherent, consistent reactor designs.

The driving idea behind BLUEPRINT is that the generation of preliminary design points in the pre-conceptual design stage should take minutes, not months.

BLUEPRINT replicates the first few stages of the design definition process, from the 0-1D systems code run to the full 3-D CAD geometry, in a tightly coupled way. All calculations performed in this loop should be fast and self-consistent with the other procedures. Other analyses are far harder to accelerate, and cannot reasonably be included in an optimisation. These are loosely coupled to the framework, and can be called upon once a satisfactory design point has been converged upon, see Fig. 3.

Ideally, for each calculation procedure in BLUEPRINT, three different levels of fidelity would exist:

- **Low:** A 0-D model or default value(s) providing the desired information rapidly, with commensurately low fidelity. For instance, consider the efficiency of the balance of plant (BOP), η_{BOP} , for a He-cooled blanket with an outlet temperature of 500° C. One could set (based on experience, or previous studies), e.g. $\eta_{BOP} = 0.33$, assuming a Brayton cycle.
- **Medium:** A higher order model, which carries out several less trivial calculations to provide the requested information with a higher fidelity for a higher computational cost. Alternatively, an “intermediate-fidelity” parameterisation based on the results of a high-

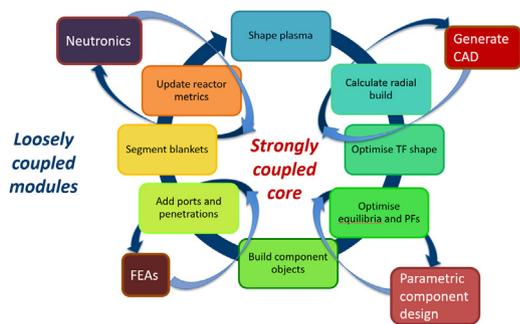


Fig. 3. Schematic of the BLUEPRINT design point definition approach. A tightly coupled set of calculations are performed consistent with one another. Loosely coupled procedures can generate other data upon request, but are not called in an optimisation loop.

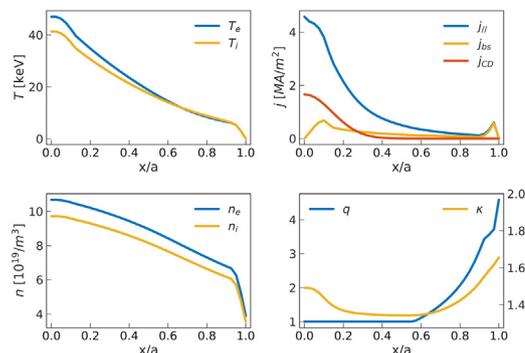


Fig. 4. Plasma profiles in Shafranov coordinates obtained using the PLASM0D code, showing (clockwise) the electron and ion temperature (T_e and T_i), the parallel, bootstrap, and driven currents ($j_{||}$, j_{bs} , and j_{CD}), the electron and ion densities (n_e and n_i), and the safety factor, q , and elongation, κ , profiles.

fidelity model can be considered. Here, one might look at previous BOP studies, and derive a parameterisation, e.g. $\eta_{BOP} = f(P_{BB}, P_{DIV}, T_{BBout})$, where P_{BB} and P_{DIV} are the thermal powers extracted from the blankets and the divertor, and T_{BBout} is the BB coolant outlet temperature.

- **High:** This would be the state-of-the-art method to obtain a certain result, with a consequently higher computational cost. A high-fidelity approach for η_{BOP} would be to automatically design the full BOP cycle for a defined set of inputs.

The different levels of fidelity and computational cost are to facilitate global optimisation; as the optimiser converges on a solution, the level of definition of some evaluation procedures can be progressively ramped up. Naturally, for a final design point, a full suite of analyses would be carried out.

Wherever possible, BLUEPRINT will aim to use open-source software. This is to enable portability, and facilitate parallel processing and deployment on high-performance computing clusters or cloud computing. There is a preference to utilise existing codes for specialist calculation procedures, where and when available.

We note that in other multi-disciplinary engineering fields where design optimisation and comparison of alternative are important, similar efforts at automated conceptual design frameworks have been elaborated. NASA, for instance, have developed the Flight Optimisation System [13], which is a system of various computer programs aimed at the conceptual design and evaluation of advanced aircraft concepts.

3.2. 0/1-D design point basis

BLUEPRINT can be initiated in two ways: either through an initial PROCESS run, with subsequent evaluations being performed on the output, or through its own internal simplified systems code optimization loop. The former is the more developed of the two, and encompasses an already well-respected and widely-used code to generate an initial 0-D reactor design.

PROCESS is called as a routine within BLUEPRINT, with access to the desired control and iteration variables, equality and inequality constraint equations, and optimisation objective. As is the custom in EU-DEMO studies, BLUEPRINT calls PROCESS with a minimisation objective for the major radius, R_0 , and prescribed values for the net electric power output, P_{elnet} , and pulse length, τ_{pulse} .

The outputs from the PROCESS code are read into memory and certain aspects are selected and used to run the following reactor design procedures. At present, BLUEPRINT relies on PROCESS for the reactor radial build, plasma current, I_p , and the field at the major radius, B_0 . The rest of the output is stored but not used, as higher-fidelity calculations in BLUEPRINT replace many of the 0-D engineering models in PROCESS.

3.3. 1-D transport solver: PLASM0D

A 1-D plasma transport and equilibrium model, PLASM0D, described in [14], is loosely coupled to BLUEPRINT. It can be used as part of the preliminary optimisation loop, or as a post-processing step on a PROCESS 0-D design point.

PLASM0D is run with a range of reactor parameters (R_0 , B_0 , I_p , but also plasma shaping parameters, physics and engineering constraints, etc.), and returns important information to BLUEPRINT. Typical outputs used in later reactor design procedures include: the fusion power, radiation corrected H-factor, normalised plasma pressure, separatrix power, bootstrap fraction, energy confinement time, etc. The normalised profiles for the plasma density, ion and electron temperatures, current, magnetic flux, etc., are also fed into BLUEPRINT (see Fig. 4).

3.4. TF coil shape optimisation procedure

The shape of the TF coil in BLUEPRINT is optimised to reduce toroidal field ripple at all positions in the separatrix to below a specified limit (usually 0.6%). The winding pack shape is then optimised according to either of two user-selected criteria: the minimisation of the winding pack (WP) perimeter, L_{TF} , or the total TF coil cage stored magnetic energy, E_{TF} . An offset surface from the outer edge of the vacuum vessel thermal shield (VVTS) is specified as an exclusion constraint. The parameterisation function of the winding pack poloidal profile is another user-selected option, with support for the Princeton-D [15,16], ITER-like D-shape (multiple circular arcs, see e.g. [17]), and a Bézier spline provided. The former two (with a total of 5 and 7 parameters, respectively) are relatively well-known in the fusion magnet community, and require little introduction. The latter is a novel, higher-order parameterisation of TF coil WP poloidal profile. It gives the user greater control, in particular when setting inequality constraints for the optimiser, enabling the fine-tuning of individual bending radii, straight length segments, etc.

The Princeton-D and the ITER-like TF coil shapes are such that they aim to minimise the in-plane bending the coils will experience due to the powerful magnetic fields they generate. However, in a superconducting tokamak, the logical design decision is to structurally and thermally connect all cryogenic masses within the torus hall. Thus, during operation, the PF coils and gravity supports apply large, varying vertical loads directly to the TF casings, negating efforts to produce a TF coil shape with low bending moments due to its own in-plane loading.

The reactor radial build, toroidal field at the major radius, B_0 , and ripple limit are inputs to the procedure, as is an exclusion zone for the TF coil—defined from an offset surface to the VVTS.

Fig. 5 shows an optimised Bézier spline TF coil, with $R_0 = 9$ m, $n_{TF} = 16$, and a ripple limit of 0.6%, and the corresponding TF ripple along the separatrix.

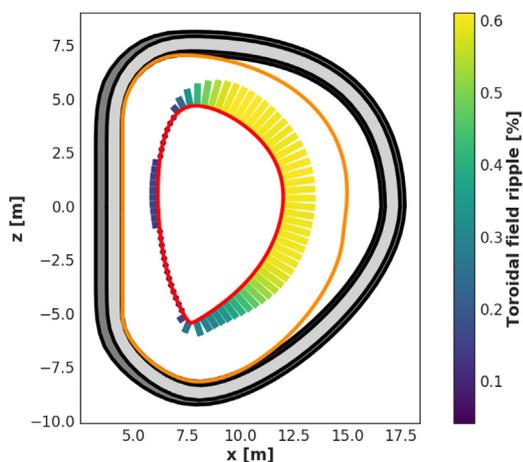


Fig. 5. Plot of the toroidal field ripple on the separatrix (shown in red) for an optimised “S” TF coil, with $n_{TF} = 16$. The TF coil exclusion zone constraint (VVTS outer surface) is plotted in orange.

Our optimisation procedure does not yet directly include structural constraints, e.g. for maximum in-plane bending. Work is under way to include a rapid, beam finite-element (FE) model into the optimiser, which will include consideration of the TF coil casing and inter-coil structures (which presently receive little attention in our procedure). However, FE analyses can be performed on the output coil designs at a later stage (see Section 3.10).

3.5. 2-D magnetic equilibrium solver: Nova

One of the key modules in BLUEPRINT is Nova: a 2-D free boundary equilibrium solver, which is packaged with associated tools for the manipulation, optimisation, and visualisation of tokamak equilibria and magnetic coils. This module will be fully described in a future publication.

The equilibrium solver is supplied with the desired separatrix shape, flux functions as internal boundary conditions, and the desired flux swing in the case of a pulsed reactor. The currents and positions of the central solenoid (CS) and PF coils are optimised to keep the root-mean-squared error of fit to the separatrix shape below a user-specified tolerance, while minimising the total absolute sum of coil currents, see Fig. 6 for a typical solution.

Following an initial optimisation of the coil positions for a given equilibrium, remote maintenance considerations for the extraction of the in-vessel components (IVCs) are usually violated (see Section 3.7). As such, an internal iteration loop can be performed within a reactor evaluation procedure, which calls Nova a second time to optimise the coil positions and currents with additional positional constraints applied to some coils.

Typically, this would mean positional constraints on the two upper most PF coils, to allow sufficient space for a vertical port through which the blanket segments may be extracted, and some positional constraints on the lower PF coils for the vacuum vessel ports for the maintenance of the divertors.

3.6. 2-D reactor build

Based on the radial build of the machine (taken either from the PROCESS run, or calculated by other means), various procedures in BLUEPRINT are used to shape the reactor components in 2-D.

3.6.1. In-vessel component segmentation

As part of the 2-D build, BLUEPRINT calls a routine to segment the in-vessel components and position the ports, such that certain criteria are met to enable a vertical extraction of the blankets. This involves a

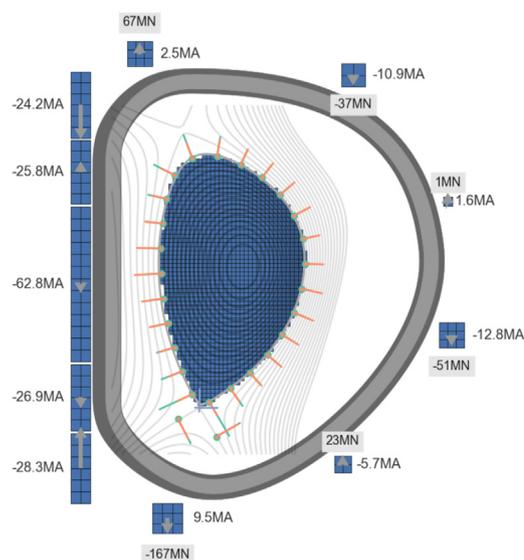


Fig. 6. Typical equilibrium solution from Nova, positioning the PF coils around the TF coils, and resulting coil currents and forces. The green and orange dots and dashes show the locations where the ψ and B values are prescribed to the optimiser.

simple representation of the blanket kinematics and some constraints for in-vessel component manipulation, and leads to a minimum size for the vertical upper port and the angled lower port. This often requires rearranging the position of the PF coils, and re-optimising the magnetic equilibria with further constraints, in order to leave sufficient space between the two uppermost PF coils to extract the blankets.

3.6.2. Plasma-facing surface shaping

Tools in the Nova module of BLUEPRINT are used to shape the blanket first wall (FW) and divertor in 2-D. Heat loads on the divertor targets and FW of a fusion reactor are often design-driving. They consist of thermal radiation, neutron radiation, and charged and neutral particle loads that exhibit high local peaking.

At present, no direct calculation of any of the aforementioned loads is performed in BLUEPRINT. However, we take an alternative approach to design the FW, using geometrical and flux surface offsets to the plasma separatrices as a proxy for the charged particle loads. Two parameters are used: (i) ψ_n , a normalised 2-D flux surface relative to the separatrix, and (ii) $\partial x \partial z$, a geometrical offset to the separatrix. A parameterised preliminary shape is fitted to a convex hull of these multiple snapshots of equilibria, which can include for example, vertically and radially displaced separatrices. Only the upper half of the flux surfaces are used in the case of a SN reactor, as divertor designs are handled separately. This is of particular importance in SN reactors, as plasma vertical stability issues and a pulsed central solenoid (CS) can lead some of the scrape-off layer power to be transported along highly elongated flux surfaces.

As flux-fitting FW shapes are likely to be overly expensive to manufacture, we discretise the preliminary FW profile obtained into flat lengths. We apply constraints to the angle between modules, α_{max} , minimum and maximum module lengths, L_{min} and L_{max} , and minimise the total number of modules. This procedure is carried out separately for inboard and outboard blanket segments, as the point of segmentation is dictated by maintenance considerations, see Fig. 7 (left).

The divertor profile is designed in a similar way around a defined X-point location, with the inboard and outboard divertor leg lengths, L_{IB} and L_{OB} , being specified by the user. The divertor target locations and flux values are constraints in the equilibrium optimiser, and the orientation of the target plates is designed such that the divertor is closed. A user input compound grazing angle of the separatrix flux surface

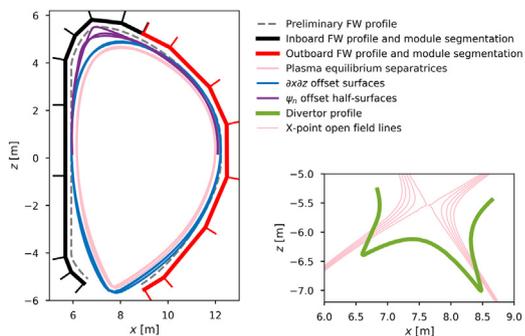


Fig. 7. Illustrative example of the algorithmic design of the plasma-facing surfaces of a SN reactor. (left) FW designed with several reference separatrices (pink), and geometric and flux surface offsets (blue and purple), where $\psi_n = 1.07$ and $\partial x \partial z = 0.225$ m. The dashed grey lines show the preliminary FW spline, and the black and red lines show the inboard and outboard modularised FW profiles. Here, $\alpha_{max} = 20^\circ$, $L_{min} = 0.75$ m, and $L_{max} = 3$ m. The inboard and outboard FW profiles would normally coincide heavily with the preliminary spline shape. They have been slightly offset here for clarity. (right) Divertor profile with $\alpha_g = 1.5^\circ$, $L_{IB} = 1.1$ m, and $L_{OB} = 1.45$ m.

relative to the target, α_g , is used to determine the angle of the target plate relative to the open field lines, see Fig. 7 (right).

The shapes of the vacuum vessel and VVTS are parameterised as multiple splines, which are sized according to the separatrix shape. A user-defined flat inboard length for these components is also specified (the design intent being to simplify manufacture and assembly). The upper and lower ports are designed to roughly accommodate the blanket and divertor maintenance, and equatorial ports of user-specified dimensions are also added. The cryostat thermal shields, cryostat vacuum vessel, and the concrete radiation shield are parameterised in a simplistic manner, see Fig. 8.

3.7. Remote maintenance calculations

Further to the initial IVC segmentation calculation, a figure of merit for the ease of maintenance is calculated: the remote maintenance technical feasibility index (RMTFI), as detailed in [18]. The RMTFI is a weighted combination of a number of technical metrics relating to the maintenance of the blanket segments, such as the torque at the lifting interface, the number of kinematics steps, etc. The RMTFI is not a fully quantitative parameter, and bears meaning only in relation to an already known and well-investigated configuration (an aspect ratio 3.1, $n_{TF} = 16$ SN tokamak, with five multi-module blanket segments per port). It gives the reactor designer (who may not necessarily be familiar with maintenance issues) quasi-instantaneous feedback on the ease of maintenance of a vertically maintained SN design. A routine in BLUEPRINT calculates the RMTFI based on the preliminary geometry information available (BB segment geometry, the location of the upper port, segment centre of gravities, etc.), see Fig. 9.

The RMTFI and the IVC segmentation procedure have been codified assuming a vertical maintenance approach, as per the current European thinking (see e.g. [19]). Alternative maintenance approaches would need to be codified within the BLUEPRINT framework for them to be compared effectively. In some cases, this may require further alternative modules to be developed if the maintenance concept fundamentally alters the machine layout. For example, a sector maintenance strategy (see e.g. [20]) would require modification or addition to the PF and TF coil modules, as well as new design parameterisations for the thermal shields, cryostat vacuum vessel, and radiation shields.

BLUEPRINT also has a model to estimate the remote maintenance durations for planned maintenance activities. This model is detailed in [21], and calculates the total shutdown durations by assigning durations to a number of elementary activities (cutting and welding of pipes, non-destructive testing of connections, etc.). Based on the sequence of

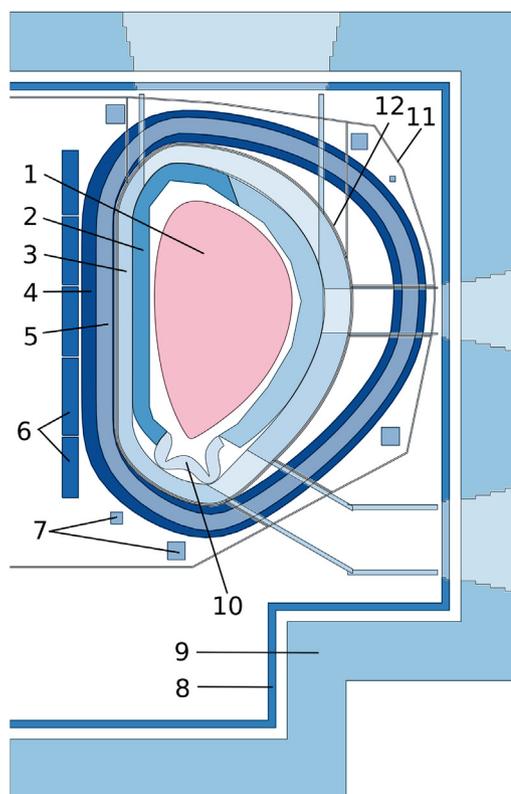


Fig. 8. Cross-section of a typical SN reactor design generated by BLUEPRINT. Legend: 1—plasma, 2—breeding blankets, 3—vacuum vessel, 4—TF coil case, 5—TF coil winding pack, 6—CS coils, 7—PF coils, 8—cryostat vacuum vessel, 9—radiation shield, 10—divertor, 11—cryostat thermal shield, 12—vacuum vessel thermal shield.

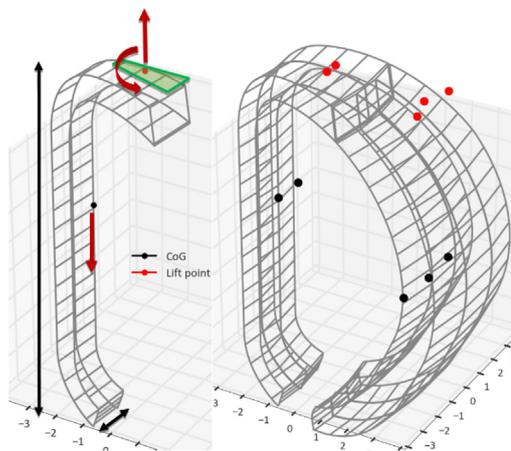


Fig. 9. Indicative plot of the important parameters required to calculate the RMTFI (segment geometries, locations of the centres of gravity, lifting locations with the port, and locations of loads and bending moments on the RM equipment).

operations required to clear all the tokamak ports and extract the in-vessel components, it carries out a critical path analysis and calculates the duration of the total intervention. For a set of reactor inputs (the maximum materials damage, the rate of accumulation of material damage, the sizes of the ports, the number of IVCs, etc.), the durations of a range of planned maintenance activities are estimated: (i) a full in-vessel component refit, (ii) the exchange of all divertor cassettes, and (iii) the exchange of different individual blanket and divertor cassettes.

3.8. Tritium breeding ratio estimation procedure

The TBR is crudely estimated from data generated by neutronic analyses for the HCPB blanket design, in which a total potential TBR estimate was calculated, and a poloidal distribution of the potential provided [22]. A simple volumetric penalty factor is calculated, which accounts for the presence of non-breeding areas (divertors, H&CD systems, high heat flux zones) and their poloidal location, poloidal depth, and toroidal extent. From the above, an estimate of the TBR is produced. This is a relatively low-fidelity approach, and at present is only available for one blanket type. It has the virtue of being computationally inexpensive and bearing some relation to actual neutronics studies.

For other blanket types, similar data from studies such as [22] would be required in order to estimate the TBR.

3.9. 3-D CAD creation

Key to more in-depth studies of reactor configuration are the 3-D geometry data (e.g. CAD artefacts). Manual creation of DEMO reactor CAD usually takes approximately one month (see Section 2.2). In BLUEPRINT, the procedure to generate a 3-D CAD model of the reactor up to the radiation shield is fully parameterised and produces a model with at least the same level of detail as is typically used in the EU-DEMO studies, see Fig. 10. The procedure itself takes approximately 10–30 seconds on a single core with an Intel® i7 processor, and a few more seconds to be loaded in 3-D CAD model-viewing freeware. This process could be accelerated by another order of magnitude (for the same level of detail) by apt use of CAD procedures and parallelisation of the various independent processes.

We concede that our cryostat vacuum vessel design is simplistic, although we also suggest that simply hanging the cryostat vacuum vessel lid from the radiation shield roof may be preferable to making a tall elliptical convex lid. Much like any other design decision for a DEMO-class reactor, alternatives (where considered) would need to be codified and selected within the BLUEPRINT framework in order to be studied.

Much like the CAD generation step in the present EU-DEMO design

process, the CAD data generated by BLUEPRINT are not intended to be highly detailed. The idea is to pass these data to dedicated studies so that further levels of detail may be added where necessary, ideally in parameterised fashion. Where the detailed conceptual design of some components affects the overall performance of the reactor significantly, our intention is to build more detailed engineering models into BLUEPRINT to account for important effects. We intend for the TF coil winding packs and the BB modules, for instance, to be treated in a more direct and detailed manner within BLUEPRINT, such that design and engineering constraints are accounted for robustly within the tightly coupled part of the design point generation procedure.

3.10. Proof of principle for coupling to engineering analysis software

For any given reactor configuration, dozens of higher-fidelity engineering analyses should be carried out for critical load cases on key reactor components. Many of these analyses are presently carried out by engineering analysts through an often non-trivial exercise in setting up the model with the correct geometry, material properties, boundary conditions, and loads. The results themselves then need to be interrogated intelligently to determine if the load cases considered are acceptable, or if adjustments to the design or inputs are required. Frequently the majority of the duration of such analyses is in establishing the model correctly. At present, much of this work is done manually, with the geometry being ported across to whichever software the analyst wishes to use. The model may need to be simplified, or detail must be added. Occasionally, the geometry is re-built altogether in a different way; constructive solid geometry neutronics models, for example.

Key to our proposal of accelerating the design cycle for fusion reactors is the capability to automate higher-fidelity analyses. In fusion engineering disciplines, these will often come in the form of neutronics or FE analyses. To demonstrate the principle of loosely coupling commercial analysis software with BLUEPRINT, we present two examples: a FE model of a TF coil, and a global neutronics model.

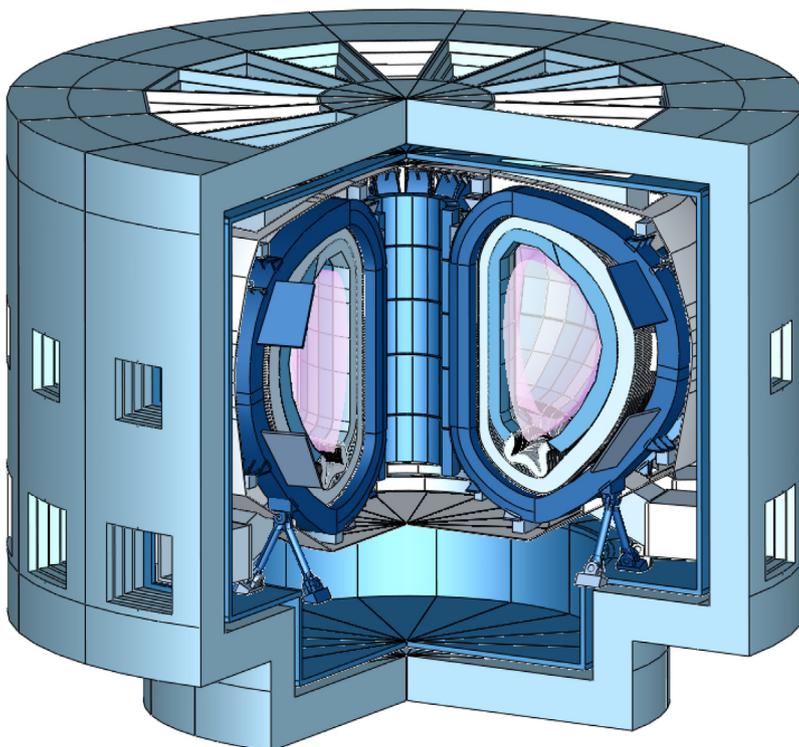


Fig. 10. Three-quarter isometric view of the 3-D CAD model of the reactor generated by BLUEPRINT in approximately 30 s, for the same design point shown in Fig. 8, with the same colour scheme. The radiation shield port plugs, cryostat vacuum vessel closure plates, and the tokamak vacuum vessel closure plates are hidden for illustrative effect.

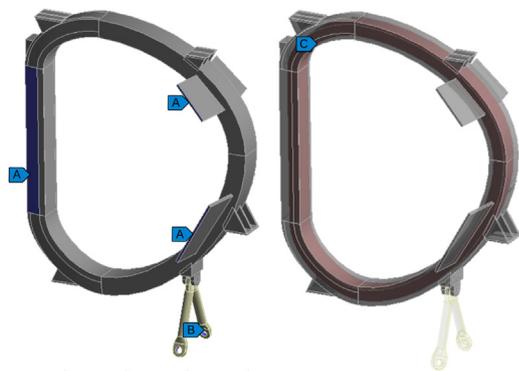


Fig. 11. Indicative ANSYS TF FE model. Showing boundary conditions: (A) cyclic symmetry; (B) frictionless support, and sliding friction joint condition between the WP and the TF coil casing (C). NOTE: The gravity supports shown here are inspired from the JT60-SA design, and have not been designed or analysed. The intercoil structures are similarly indicative; indeed this coil would likely fail under the combined loads in its present form.

3.10.1. Finite element model: proof of principle

Here we use ANSYS [23], although the principle of coupling to other commercial or open-source FE software will be similar. In practice, it is likely that there will be a variety of syntactical differences and implementation issues. We note however that this is a software engineering problem, and once solved, would result in significant savings of cost and time.

The geometry is loaded as a STEP file format, and assigned material properties. The PF vertical forces and the gravity support loads are applied as point loads. The TF coil bursting and toppling forces are applied to the mesh nodes in the WP using an ANSYS macro. Boundary conditions can be applied to named surfaces specified during the creation of the STEP file. As such, the geometry, boundary conditions, and loads can be loaded directly into an analysis package, see Fig. 11.

While this procedure is not yet fully automated, based on our preliminary efforts it is entirely conceivable that it could be. Moving to a freely available and more script-based FE software package is our next step in this direction.

In principle, multiple such higher fidelity analyses could be loosely coupled to BLUEPRINT, and run once an attractive design point has been identified. With enough high-fidelity data points, it could be possible to develop intermediate-fidelity models; where values are predicted from data rather than explicitly calculated.

3.10.2. Global neutronics model: proof of principle

Using the same geometry objects created in BLUEPRINT, we can also export parts in a stereolithography file format (STL). STLs are surface meshes and, while less computationally convenient than the constructive solid geometry (CSG) file format—often used in neutronics studies—STL files present a number of advantages. Crucially, no transformation of CAD geometries into CSG objects is required, and volumetric surface meshes can be created directly from a CAD object. Care must be taken to ensure that the surface meshes are coherent, “watertight”, and of reasonable quality. This is crucial step in the process, and in the past has been the source of much delay to neutronics calculations, owing to time spent manually defeating and/or fixing geometry objects. At present, the geometries generated in BLUEPRINT are fairly simple, but with increasing levels of detail, geometry post-processing workflows will become necessary. Many neutronicists have tackled such problems in the past, and various CAD mesh processing workflows have been developed (see e.g. [24]), which we would seek to use.

We initialise all of the major reactor component 3-D CAD data into a Serpent II [25] model as STL files, define a standard parameterised D-T neutron source, as described in [26] and used commonly in EU-DEMO

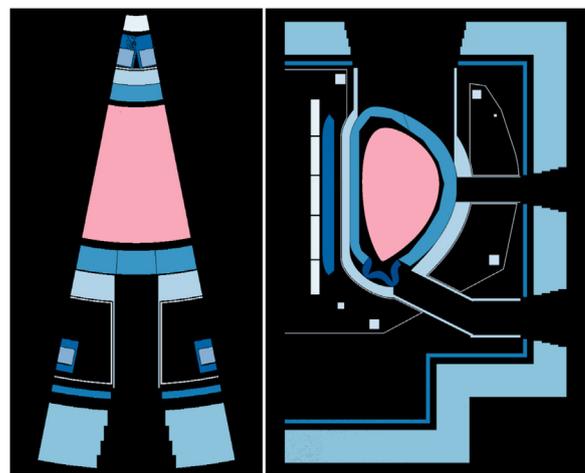


Fig. 12. Indicative Serpent II neutronics modelled initialised from the same reactor object as shown in Figs. 8 and 10. (left) Plan view of a sector slice at $z = 0$. (right) Poloidal cross-section in between two TF coils.

neutronics studies, and assign typical homogenised material properties (e.g. as specified in [27]) to all the components, see Fig. 12. Unlike in the example FE model, no complicated boundary conditions are required, and there is only one load (the plasma).

Were we to add more heterogeneous representations of the blankets, divertor, and vacuum vessel, this model (when run with an adequate number of neutron histories) would readily return the TBR, neutron and gamma power deposition integrals, neutron fluxes, and material damage rates in key areas.

3.11. Lifecycle module

The plant lifecycle is important for a variety of reasons, especially when considering short-pulsed reactor operation. Typical EU-DEMO studies assume a very low load factor, as a first-of-a-kind demonstrational fusion power plant with many complex sub-systems is unlikely to achieve high availability. It is important to probe the consequences of this from the perspectives of the tritium fuel cycle, as low load factors early in the DEMO operational phase can drive the amount of tritium start-up inventory required. We also note that fatigue issues deserve attention, as may systems which are unaccustomed to receiving sporadic, irregular loads (e.g. BOP systems).

To address the above issues, we have built a lifecycle generation procedure, which generates fusion power loading patterns over the lifetime of the plant, which are partly based on the expected operational schedule of the reactor, and partly randomised - mimicking the effects of unpredictable sub-system failures and outages.

The lifecycle objects are structured around when scheduled replacements of in-vessel components take place. These events are calculated based on defined in-vessel component lifetimes (expressed in terms of displacements per atom or total neutron fluence above an energy threshold) and neutron fluxes estimated crudely from previous studies and defined shield thicknesses. In future, these values will be obtained directly from loosely coupled neutronics analyses within BLUEPRINT.

The maintenance shutdown durations are taken from the RM model explained in Section 3.7. Based on a global load factor target over the operational life of the plant specified by the user, the procedure then distributes operational availabilities to the different operational phases (between scheduled maintenance operations). The distribution takes into account a sigmoid-type learning curve, accounting for gains in operational experience over the life of the plant which compete with the effects of infant mortality and wear-out failures. Then, within each phase, the total outage is distributed between the pulses, using a log-

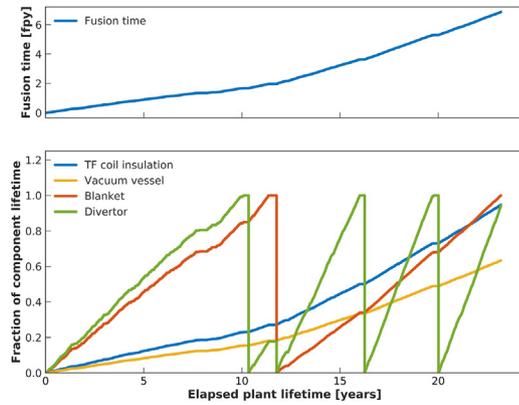


Fig. 13. Example of a DEMO-class reactor lifecycle generated by BLUEPRINT. Above the evolution of the reactor fpy over its life is shown, and below is the build-up of damage in key reactor components throughout the course of operation of the reactor.

normal distribution for the duration of the inter-pulse durations, with a defined integral. Thus the inter-pulse durations are made to vary between a specified (or calculated) minimum (usually the CS recharge time, e.g. ~ 600 s) and theoretically $+\infty$, although this is limited by the integral of the distribution and in practice can reach up to several months. More detail on this procedure can be found in [28].

For the purposes of illustration, Fig. 13 shows the fusion time in terms of full-power years (fpy) and the elapsed plant lifetime in calendar years. The changing slope is dictated by the learning curve assumed, which modifies the average load factor in each operational phase. Also shown is the accumulated damage to the blankets, divertors, vacuum vessel, and TF coil insulation over the course of the life of the plant. This is shown in terms of the fraction of component lifetime, with respect to its defined limits (e.g. 20 dpa for the starter blanket). The blankets are replaced once and the divertor on several occasions, as per the current EU-DEMO thinking.

3.12. Fuel cycle module

The fusion power loading over the life of the plant can be used in a variety of ways. Here we briefly introduce one important application: a dynamic tritium fuel cycle model. As key authors have made clear in their works [29,30], understanding the fuel cycle is vital to understanding tritium self-sufficiency, and what is needed in terms of the performance of the blankets (TBR), plasma (burn-up fraction, f_b), and TFV systems (Direct Internal Recycling [31] factor, f_{DIR} , residence time, etc.), and overall reactor performance (global load factor, A_{glob}). With a similar intent to these authors, we have incorporated a new fuel cycle model in BLUEPRINT, which is described in more detail in [28].

For a given reactor configuration and a set of performance parameters for the tritium, fuelling, and vacuum systems, the fuel cycle model is capable of estimating the tritium start-up inventory, $m_{T_{start}}$, and the reactor doubling time, t_d . The model is run several hundred times for a given reactor and tritium fuelling and vacuum (TFV) system configuration using different randomly generated timelines, for which the worst-case values of $m_{T_{start}}$ and t_d are returned. Fig. 14 shows a typical output from the fuel cycle module, showing the evolution of the tritium inventory throughout the life of the plant for one timeline, accounting for natural decay, sequestration in various components, and production in the blankets.

3.13. Reactor power balance

BLUEPRINT calculates the reactor power balance in a similar way to existing systems codes, with presently only simple representations of the BOP secondary cycle (where the same parameterisations for the

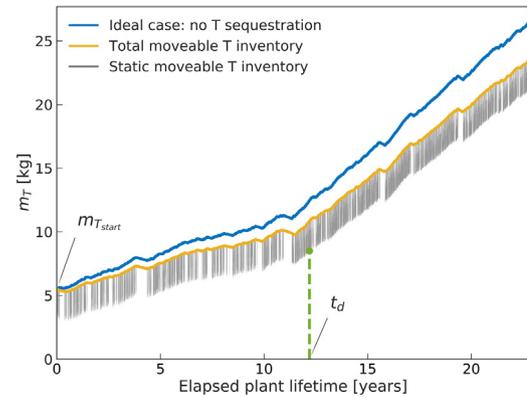


Fig. 14. Tritium inventory evolution over the life of the plant. The blue shows an indicative case where no tritium loss terms for sequestration in the IVCs or TFV systems are included. The grey lines show the fluctuations of the tritium throughout every pulse over the history of the machine: as it is taken from the stores, used to fuel the plasma, and eventually returned to the stores by various TFV systems. In this example: $TBR = 1.05$, $f_b = 1.5\%$, $A_{glob} = 0.3$, and $f_{DIR} = 0.8$. For an exhaustive description, please see [28].

cycle efficiency as in PROCESS are used). Similarly to the estimation of the TBR, volumetric ratios of the in-vessel components are used to estimate the neutron heating in the various plasma-facing components (including volumes allocated to heating and current drive systems), and a correction factor is used to represent the neutron heating in the VV and TF coils.

Fig. 15 shows the full thermal and electrical power flows through the fusion reactor model. The pumping power required to cool the blankets and divertors is estimated as a fraction of the incident heating power or a pressure drop based on previous studies. In future, these models will be replaced with higher-fidelity ones which bear some relation to the actual designs of the IVCs. The remainder of the parasitic loads (cryoplant, tritium plant, magnet electrical power) are also estimated crudely with 0-D models, similar to those used in existing systems codes.

4. Comparison of the EU-DEMO design process and BLUEPRINT

We do not consider the time to run the neutronics model in our comparison between the present approach and the BLUEPRINT runtime. This is because our neutronics model is not detailed enough for a fair comparison, and because it stands to reason that we would not run $\sim 10^9$ neutron histories on a single core of a desktop computer.

Instead we compare the time it takes to achieve all steps in the design process up to the generation of the CAD data (see Fig. 2). For the current design process this takes approximately 1.5 months. For typical SN reactor runs on a single core with an Intel i-7 processor, from the PROCESS run all the way through to and including the generation of the 3-D CAD data, BLUEPRINT runtimes vary between two and four minutes.

The level of definition is harder to compare; however, we consider that overall BLUEPRINT matches or exceeds typical levels of definition of recent EU-DEMO design point iterations. In some aspects of importance, the tightly coupled analyses in BLUEPRINT go above and beyond the normal DEMO design process; e.g. checking and measuring remote maintenance compatibility and feasibility (albeit in a preliminary fashion), and FW shape optimisations. In other areas (e.g. PROCESS run), BLUEPRINT is identical to the standard DEMO design procedure. All things accounted for, we contend that BLUEPRINT generates design points of the same order of depth and fidelity, but four orders of magnitude faster than present methods.

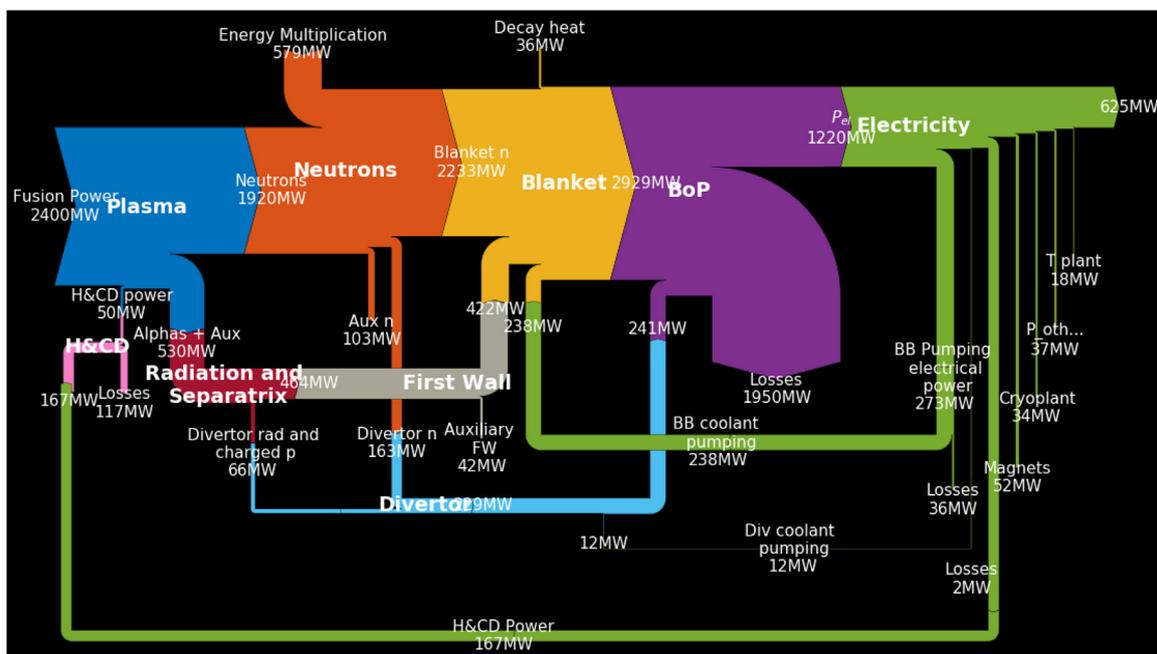


Fig. 15. Sankey diagram of an indicative pulsed DEMO power plant with $P_{fus} = 2.4$ GW, HCPB blankets, and NBI current drive.

5. Discussion

The fusion reactors that the fusion community designs today will not be those built in the 2040s. There is much to discover and learn from ongoing design studies, R&D activities, and of course, the operation of ITER. Experience from ITER illustrates the importance of even a relatively detailed design remaining flexible throughout the engineering design activities, and responsive to disruptive technologies, fortunate and unfortunate developments in R&D, and unforeseeable modifications to the design constraints.

As a framework for combining, automating, and streamlining typical conceptual design and analysis studies for future fusion reactors, BLUEPRINT can help to steer the design and guide R&D. Where it is important to investigate reactor performance over a certain parameter range (e.g. for the plasma elongation or aspect ratio), BLUEPRINT would be well suited to investigate a far larger number of data points much faster than present methods. Indeed, a subject of further work will be to revisit and expand upon the aspect ratio study carried out in [32] (over three data points), to demonstrate the application of BLUEPRINT and compare outputs against existing results.

Furthermore, preliminary efforts have shown that in principle alternative magnetic geometries (another EU-DEMO key design issue, see e.g. [33]) can also be handled by BLUEPRINT. This said, certain activities would be very challenging to automate or proceduralise. Engineering integration issues in particular would be difficult to handle algorithmically; reactor designers are needed to identify and resolve such problems, which typically require a broad range of expertise and difficult design decisions to be made. Indeed, some key design issues in the EU-DEMO programme fall into this category, such as the blanket maintenance design issue (see [9]). Similarly, there is no substitute for technology R&D activities; they are vital to inform reactor design codes.

For the studies which cannot be parameterised or automated, we propose instead that BLUEPRINT be used as a central repository for information for reactor design points. To carry out such a study, engineers and physicists should be able to extract all of their requirements, space allocations, load cases, fields, etc., from a reactor design object, which should be truly consistent with the values and assumptions being used in other studies. To enable loosely coupled high fidelity and/or multi-physics studies to be carried out, it will be of importance in future to utilise standardised mesh and data structures, for

passing e.g. a temperature field back and forth between a neutron transport simulation, computational fluid dynamics simulation, and thermo-mechanical FE analyses.

A key objective for BLUEPRINT is to bridge the gap between instantaneous 0-D system code studies and the detailed multi-annual studies. This will pave the way for intermediate-fidelity reduced models to be constructed, enabling complex engineering analyses to be incorporated within optimisation loops which would otherwise be prohibitively expensive.

A faster design and analysis procedure will hopefully enable us to learn from our design experiments, innovate more effectively, and explore a wider design space. The ability to handle a broad range of design options for DEMO-class reactors will be crucial to properly searching the design space, and ultimately to ruling out poor reactor or sub-system concepts once and for all.

Today, fusion reactors across the world are designed by different teams with different ideas, using different assumptions, see [2,5,34–36]. Despite having more or less the same objectives, these paper reactors are largely incomparable.

In the medium term, our aim is to be able to model at least a few different reactor concepts (focussing on those selected for study in the EU, see e.g. [33]), using a common understanding of physics and engineering limits and constraints. Only then can we truly say which is better based on what we know today. Only then can we down-select the various technologies and designs which we consider.

6. Further work

Here we only intend to demonstrate the first steps towards an advanced systems code that can significantly speed up the conceptual definition of a reactor design point and pave the way for rapid, meaningful parameter scans, trade-off studies and paper reactor comparisons. Much work remains to be done:

- Tightly coupled calculations (inside the optimisation loop)
 - 2-D beam FE structural analyses for the magnetic cage, incorporating out-of-plane loads to the analysis and optimisation of the TF coil shape and PF coil currents and positions
 - 2-D method of characteristics neutronics model for intermediate-fidelity TBR and power deposition estimates

- Higher-fidelity TF coil winding pack models
- Unsteady equilibria and plasma vertical stability calculations
- ...
- Loosely coupled calculations
 - Full, representative global neutronics model with high-fidelity TBR and power deposition estimates
 - 3-D FE thermal and structural analyses for the full magnetic cage
 - 3-D vacuum pumping conduction analyses
 - Higher-fidelity BOP models
 - ...
- Alternative sub-system concepts
 - Water-cooled blankets
 - Pancake and layer-wound TF coil winding packs
 - ...
- Alternative reactor designs
 - Double-null
 - Snowflake
 - Super-X
 - Spherical tokamak
 - ...

Of course, many more analyses and capabilities could be imagined. It will be of importance to understand which are required, and to what degree of fidelity, for which configurations. Despite having developed some of the above aspects further than we have shown here, we do not pretend to be able to deliver all of them unassisted. Ideally, routines for such analyses would be written by collaborators expert in the relevant domains. To this aim, we invite and welcome collaborations, voluntary contributions of modules, models, procedures, or otherwise, provided that they make sense for the design of tokamaks.

7. Conclusions

We have outlined some of our thoughts regarding the present approach taken for the design of EU-DEMO reactor design points, highlighting some weaknesses and the issues that arise.

To demonstrate an improved approach, a novel reactor design code, BLUEPRINT, has been created, and its present capabilities are introduced here in a preliminary fashion. We show that the reactor design point definition process can be automated and accelerated by four orders of magnitude—from months to minutes.

Acknowledgements

We would like to thank Dr E. Fable for giving us access to PLASMOD and his help using it, Dr H. Lux for her help with PROCESS, I. Farquhar for his model on remote maintenance shutdown durations, Dr J. Shimwell and Dr A. Davis for their help and guidance with neutronics models, E. Jensen for her English, and Dr G. Federici, Dr G. Parks and Dr E. Surrey for their mentorship, support, wise words and guidance.

This work has been funded by the RCUK Energy Programme [Grant Number EP/I501045]. To obtain further information on the data and models underlying this paper please contact PublicationsManager@cfe.ac.uk.

References

- [1] F. Romanelli, A roadmap to the realization of fusion energy, E2C 2013. 3. European Energy Conference (2012).
- [2] G. Federici, C. Bachmann, L. Barucca, W. Biel, L. Boccaccini, R. Brown, C. Bustreo, S. Ciattaglia, F. Cismondi, M. Coleman, V. Corato, C. Day, E. Diegele, U. Fischer, T. Franke, C. Gliss, A. Ibarra, R. Kembleton, A. Loving, F. Maviglia, B. Meszaros, G. Pintsuk, N. Taylor, M.Q. Tran, C. Vorpahl, R. Wenninger, J.H. You, DEMO design activity in Europe: progress and updates, *Fusion Eng. Des.* (2018).
- [3] G. Federici, R. Kemp, D. Ward, C. Bachmann, T. Franke, S. Gonzalez, C. Lowry, M. Gadomska, J. Harman, B. Meszaros, C. Morlock, F. Romanelli, R. Wenninger, Overview of EU DEMO design and R&D activities, *Fusion Eng. Des.* 89 (2014) 882–889.
- [4] G. Federici, C. Bachmann, W. Biel, L. Boccaccini, F. Cismondi, S. Ciattaglia, M. Coleman, C. Day, E. Diegele, T. Franke, M. Grattarola, H. Hurlzmeier, A. Ibarra, A. Loving, F. Maviglia, B. Meszaros, C. Morlock, M. Rieth, M. Shannon, N. Taylor, M.Q. Tran, J.H. You, R. Wenninger, L. Zani, Overview of the design approach and prioritization of R&D activities towards an EU DEMO, *Fusion Eng. Des.* 109–111 (2016) 1464–1474.
- [5] Y. Wan, J. Li, Y. Liu, X. Wang, V. Chan, C. Chen, X. Duan, P. Fu, K. Xiang Gao, S. Feng, Y. Liu, P. Song, B. Weng, F. Wan, H. Wan, Wang, M. Songtao Wu, Q. Ye, G. Yang, G. Zheng, Q. Zhuang, C. Li, team, Overview of the present progress and activities on the CFETR, *Nucl. Fusion* 57 (10) (2017) 102009.
- [6] G. Giruzzi, J.F. Artaud, M. Baruzzo, T. Bolzonella, E. Fable, L. Garzotti, I. Ivanova-Stanik, R. Kemp, D.B. King, M. Schneider, R. Stankiewicz, W. Stpniewski, P. Vincenzi, D. Ward, R. Zagórski, Modelling of pulsed and steady-state DEMO scenarios, *Nucl. Fusion* 55 (7) (2015) 073002.
- [7] F.P. Orsitto, R. Villari, F. Moro, T.N. Todd, S. Lilley, I. Jenkins, R. Felton, W. Biel, A. Silva, M. Scholz, J. Rzakdziewicz, I. Duran, M. Tardocchi, G. Gorini, C. Morlock, G. Federici, A. Litnovsky, Diagnostics and control for the steady state and pulsed tokamak DEMO, *Nucl. Fusion* 56 (2) (2016) 026009.
- [8] P. Bunting, D. Iglesias, Assessment and Divertors Integration Studies in a DN Configuration, Tech. Rep. EFDA_D_2QSGV, UKAEA, 2018.
- [9] G. Federici, M. Siccino, R. Kembleton, Overview of the DEMO design staged approach in Europe, *Fusion Eng. Des.* (2019) (in press).
- [10] M. Kovari, R. Kemp, H. Lux, P. Knight, J. Morris, D.J. Ward, “PROCESS”: a systems code for fusion power plants-Part 1: Physics, *Fusion Eng. Des.* 89 (2014) 3054–3069.
- [11] M. Kovari, F. Fox, C. Harrington, R. Kembleton, P. Knight, H. Lux, J. Morris, “PROCESS”: a systems code for fusion power plants – Part 2: Engineering, *Fusion Eng. Des.* 104 (2016) 9–20.
- [12] C. Reux, L.D. Gallo, F. Imbeaux, J.-F. Artaud, P. Bernardi, J. Bucalossi, G. Ciruolo, J.-L. Duchateau, C. Fausser, D. Galassi, P. Hertout, J.-C. Jaboulay, A. Li-Puma, B. Saoutic, L. Zani, I.-T. Contributors, DEMO reactor design using the new modular system code SYCOMORE, *Nucl. Fusion* 55 (7) (2015) 073011.
- [13] National Aeronautics and Space Administration, Flight Optimization System (FLOPS) Software v.9, (2018).
- [14] E. Fable, C. Angioni, M. Siccino, H. Zohm, Plasma physics for fusion reactor system codes: framework and model code, *Fusion Eng. Des.* 130 (2018) 131–136.
- [15] J. File, R.G. Mills, G.V. Sheffield, Large superconducting magnet designs for fusion reactors, *IEEE Trans. Nucl. Sci.* 18 (1971) 277–282.
- [16] S.L. Gralnick, F.H. Tenney, Analytic solutions for constant-tension coil shapes, *J. Appl. Phys.* 47 (1976) 2710–2715.
- [17] F. Savary, R. Gallix, J. Knaster, N. Mitchell, K. Seo, The toroidal field coils for the ITER project, *IEEE Trans. Appl. Supercond.* 22 (2012) 4200904.
- [18] D. Wolff, M. Coleman, D. Cooper, O. Crofts, J. Horne, J. Keep, A. Loving, The impact on remote maintenance of varying the aspect ratio and number of TF coils for DEMO, *Fusion Eng. Des.* 124 (2017) 553–557.
- [19] O. Crofts, A. Loving, D. Iglesias, M. Coleman, M. Siuko, M. Mittwollen, V. Queral, A. Vale, E. Villedieu, Overview of progress on the European DEMO remote maintenance strategy, *Fusion Eng. Des.* 109–111 (2016) 1392–1398.
- [20] L.M. Waganer, ARIES-AT maintenance system definition and analysis, *Fusion Eng. Des.* 80 (2006) 161–180.
- [21] I. Farquhar, Maintenance Duration Estimator Tool, Tech. Rep. EFDA_D_2NBXBQ, RACE, UKAEA, 2017.
- [22] P. Pereslavtsev, C. Bachmann, U. Fischer, Neutronic analyses of design issues affecting the tritium breeding performance in different DEMO blanket concepts, *Fusion Eng. Des.* 109–111 (2016) 1207–1211.
- [23] ANSYS Academic Research, 2017.
- [24] A. Davis, J. Barzilla, A. Ferrari, K.T. Lee, V. Vlachoudis, P.P.H. Wilson, FluDAG: a CAD based tool for high energy physics, *Nucl. Instr. Methods Phys. Res. Sect. A: Accel. Spectrom. Detect. Assoc. Equip.* 915 (2019) 65–74.
- [25] J. Leppänen, M. Pusa, T. Viitanen, V. Valtavirta, T. Kaltiaisenaho, The Serpent Monte Carlo code: status, development and applications in 2013, *Ann. Nucl. Energy* 82 (2015) 142–150.
- [26] C. Fausser, A.L. Puma, F. Gabriel, R. Villari, Tokamak D-T neutron source models for different plasma physics confinement modes, *Fusion Eng. Des.* 87 (2012) 787–792.
- [27] U. Fischer, Guidelines for Neutronics Analyses, Tech. Rep. EFDA_D_2L8TR9, Karlsruhe Institute for Technology, 2016.
- [28] M. Coleman, Y. H”orstensmeyer, F. Cismondi, DEMO dynamic tritium fuel cycle model and plant high level requirements, *Fusion Eng. Des.* (2018) (submitted for publication).
- [29] M.A. Abdou, E.L. Vold, C.Y. Gung, M.Z. Youssef, K. Shin, Deuterium-tritium fuel self-sufficiency in fusion reactors, *Fusion Technol.* 9 (1986) 250–285.
- [30] W. Kuan, M.A. Abdou, A new approach for assessing the required tritium breeding ratio and startup inventory in future fusion reactors, *Fusion Technol.* 35 (1999) 309–353.
- [31] C. Day, T. Giegerich, The Direct Internal Recycling concept to simplify the fuel cycle of a fusion power plant, *Fusion Eng. Des.* 88 (2013) 616–620.
- [32] R. Wenninger, R. Kembleton, C. Bachmann, W. Biel, T. Bolzonella, S. Ciattaglia, F. Cismondi, M. Coleman, A.J.H. Donné, T. Eich, E. Fable, G. Federici, T. Franke, H. Lux, F. Maviglia, B. Meszaros, T. Pütterich, S. Saarema, A. Snickers, F. Villone, P. Vincenzi, D. Wolff, H. Zohm, The physics and technology basis entering European system code studies for DEMO, *Nucl. Fusion* 57 (1) (2017) 016011.
- [33] R. Kemp, R. Wenninger, G. Federici, H. Reimerdes, R. Ambrosino, R. Brown, M. Coleman, Exploring a broad spectrum of design options for DEMO, *Fusion Eng. Des.* (2018).
- [34] K. Tobita, N. Asakura, R. Hiwatari, Y. Someya, H. Utoh, K. Katayama, A. Nishimura,

- Y. Sakamoto, Y. Homma, H. Kudo, Y. Miyoshi, M. Nakamura, S. Tokunaga, A. Aoki, the Joint Special Design Team for Fusion DEMO, Design strategy and recent design activity on Japan's DEMO, *Fusion Sci. Technol.* (2017) 1–9.
- [35] K. Kim, K. Im, H. Kim, S. Oh, J. Park, S. Kwon, Y. Lee, J. Yeom, C. Lee, G.-S. Lee, G. Neilson, C. Kessel, T. Brown, P. Titus, D. Mikkelsen, Y. Zhai, Design concept of K-DEMO for near-term implementation, *Nucl. Fusion* 55 (2015) 053027.
- [36] B.N. Sorbom, J. Ball, T.R. Palmer, F.J. Mangiarotti, J.M. Sierchio, P. Bonoli, C. Kasten, D.A. Sutherland, H.S. Barnard, C.B. Haakonsen, J. Goh, C. Sung, D.G. Whyte, ARC: a compact, high-field, fusion nuclear science facility and demonstration power plant with demountable magnets, *Fusion Eng. Des.* 100 (2015) 378–405.