

This is a repository copy of *Towards a fusion power plant : Integration of physics and technology*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/188061/>

Version: Published Version

Article:

Morris, A. W., Akers, R. J., Cox, M. et al. (5 more authors) (2022) Towards a fusion power plant : Integration of physics and technology. Plasma Physics and Controlled Fusion. 064002. ISSN 1361-6587

<https://doi.org/10.1088/1361-6587/ac6694>

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

PAPER • OPEN ACCESS

Towards a fusion power plant: integration of physics and technology

To cite this article: A W Morris *et al* 2022 *Plasma Phys. Control. Fusion* **64** 064002

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

You may also like

- [Advance in the conceptual design of the European DEMO magnet system](#)
K Sedlak, V A Anvar, N Bagrets et al.
- [Advances in the physics basis for the European DEMO design](#)
R. Wenninger, F. Arbeiter, J. Aubert et al.
- [Proposal of the confinement strategy of radioactive and hazardous materials for the European DEMO](#)
X.Z. Jin, D. Carloni, R. Stieglitz et al.




physicsworld

AI in medical physics week

20–24 June 2022

Join live presentations from leading experts
in the field of AI in medical physics.

physicsworld.com/medical-physics



Towards a fusion power plant: integration of physics and technology

A W Morris^{1,*} , R J Akers¹, M Cox¹, F Militello¹ , E Surrey¹ , C W Waldon¹,
H R Wilson²  and H Zohm³

¹ UKAEA-CCFE, Culham Science Centre, Abingdon, Oxon OX14 3DB, United Kingdom

² York Plasma Institute, Department of Physics, University of York, Heslington, York YO10 5DD, United Kingdom

³ Max-Planck-Institut für Plasmaphysik, Boltzmannstr. 2, 85748 Garching, Germany

E-mail: william.morris@ukaea.uk

Received 24 December 2021, revised 31 March 2022

Accepted for publication 12 April 2022

Published 5 May 2022



CrossMark

Abstract

A fusion power plant can only exist with physics and technology acting in synchrony, over space (angstroms to tens of metres) and time (femtoseconds to decades). Recent experience with the European DEMO programme has shown how important it is to start integration early, yet go deep enough to uncover the integration impact, favourable and unfavourable, of the detailed physical and technological characteristics. There are some initially surprising interactions, for example, the fusion power density links the properties of materials in the components to the approaches to waste and remote maintenance in the context of a rigorous safety and environment regime. In this brief tour of a power plant based on a tokamak we outline the major interfaces between plasma physics and technology and engineering considering examples from the European DEMO (exhaust power handling, tritium management and plasma scenarios) with an eye on other concepts. We see how attempting integrated solutions can lead to discoveries and ways to ease interfaces despite the deep coupling of the many aspects of a tokamak plant. A power plant's plasma, materials and components will be in new parameter spaces with new mechanisms and combinations; the design will therefore be based to a significant extent on sophisticated physics and engineering models making substantial extrapolations. There are however gaps in understanding as well as data—together these are termed 'uncertainties'. Early integration in depth therefore represents a conceptual, intellectual and practical challenge, a challenge sharpened by the time pressure imposed by the global need for low carbon energy supplies such as fusion. There is an opportunity (and need) to use emerging transformational advances in computational algorithms and hardware to integrate and advance, despite the 'uncertainties' and limited experimental data. We use examples to explore how an integrated approach has the potential to lead to consistent designs that could also be resilient to the residual uncertainties. The paper may stimulate some new thinking as fusion moves to the design of complete power plants alongside an evolving and maturing research programme.

* 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, DEMO, integration, epistemic uncertainty

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

1. Introduction and context—why does integration matter

It is increasingly recognised that effective fusion power plant design depends on early integration of the physics of plasma and materials with technology and engineering, over a wide range of space and time scales and a wide spectrum of interacting physics and engineering mechanisms. As design integration is explored it becomes clear that for a fusion plant with its many ‘moving parts’ there are strong interactions between elements that might initially have been considered independent or handled sequentially. For example, details of the plasma scenario can significantly affect the overall power economy of the plant and link with the maintenance and waste strategies. Integration also extends to cost (capital, operational, lifetime); timelines and development of the supply chains; safety and environmental considerations (and thus regulation); all of which can reach far back into the design and concept choices. It is helpful and important to design backwards from the end, i.e. consider the desired end states in detail. On the other hand, too tight an integration too early can risk stifling innovation and even blocking solutions, as well as potentially adding ‘viscosity’ to the design process, so a balance has to be struck. Moving from current experiments to fusion power plants involves extrapolations in physics, technology and engineering that involve significant uncertainties. While further experimental and theoretical research and development will reduce some of these uncertainties, they cannot be avoided and have to be managed. This paper uses a variety of examples to explore ways in which an integrated approach to physics and technological design choices has the potential to lead to designs that could be resilient to the residual uncertainties.

There are already examples where considering integration potentially allows ideas in one area to reduce the challenge/open solutions in another: integration needn’t only make problems more constrained and challenging, as will be seen later in the areas of tritium, exhaust and perhaps even tokamak disruptions. The concept of a ‘detached’ divertor plasma not only eases the challenge of the plasma facing components (PFCs), it could essentially eliminate it in some places; accepting a fixed ratio of D:T in the plasma could radically reduce the scale of the tritium plant; a really effective disruption mitigation system might ease the demands on the vacuum vessel and allow blanket designs to be less constrained. The advent of high temperature superconductors and stronger steels could open up avenues for higher field devices operating further from stability limits, devices that could also be more compact.

The goal of fusion research to provide a large contribution to the global energy supply on a timescale relevant to sustained mitigation of climate change brings time pressure: while fusion power is only expected to enter after the immediate climate change impacts have been mitigated, time (and cost) still

appear to prohibit more traditional approaches to engineering and technology development with designs and decisions based on comprehensive experimental data and incremental prototypes. Large advances are needed from each stage with large extrapolations to the next—consider the step from JET to ITER, and from ITER to an ITER-like DEMO (DEMONstration fusion device) and there are more extreme examples, including plans of some private fusion enterprises. One consequence is that steps have to be conducted in parallel; this is already apparent in Europe where a DEMO [1–3] is being designed before high performance operational data from ITER is available even though the DEMO strategy is predicated upon ITER success.

These extrapolations and other aspects described later lead to many uncertainties which need to be handled. ‘Uncertainty’ is used here as a technical term describing scatter or gaps in experimental and modelling data, imprecisely known initial and boundary conditions and disturbances, gaps in understanding—it should not be equated with doubt or worry. As a result identification, quantification and management of uncertainty is becoming an intrinsic feature of design and will be a theme throughout this paper; some oversight of the overall uncertainty and its management will also be important, and represents a different challenge: how does one work out the overall plant performance uncertainty due to the presence of knowledge gaps in different areas as well as the more traditional uncertainties?

A further consequence of the ambitious strategies and schedules is that the design of a power plant will almost certainly be based on theoretical models with some (limited) support from experimental data. This reliance on models can bring many advantages, especially in handling uncertainties—but it needs the models to be adequate and timely.

There are many elements of a fusion plant which need to be brought together: all the technical aspects (which have different roles and behaviour over time) together with a wide range of capabilities in the community and in industry. Building a supply chain to provide the materials, components, skilled and experienced STEM⁴ individuals and then construct individual plants followed by fleets is a complex task that will take many years/decades as well as imagination and investment. Indeed, the ability to develop a supply chain can influence or even determine design choices—for example some materials (e.g. sufficiently enriched lithium, affordable low activation radiation-resilient steels), as well as manufacturing capability (e.g. for very large components such as toroidal field (TF) coils and vessel sections). This paper focuses on a subset of technical integration tasks, but these wider aspects should be held in mind, as should safety and environment acceptability (accidents, waste) and cost (capital, operational,

⁴ Science, Technology, Engineering and Mathematics.

lifetime)—designing with cost in mind, or its proxies such as complexity, tolerance, manufacturing challenge, is essential.

Taking consideration of this (evolving) background, the paper lays out a number of questions and theses on integrated design and gives some examples and ideas of how it might be achieved. While similar considerations apply to all approaches to fusion, e.g. inertial confinement, magnetised target fusion and magnetic confinement in various configurations (primarily tokamaks and stellarators), we focus on tokamaks (although many aspects transfer readily to stellarators).

There are other reasons to address science and technology integration early—it drives consistency, and provides a strong focus for the goal driven parts of fusion R&D programmes in parallel with blue skies exploratory R&D. There are examples where it stimulates and opens up improved solutions, or shows how to resolve challenges, as well as cases where it reveals increased challenges needing innovation, or identifies conflicts. Developing effective (and fast) integration tools will assist rapid redesign later, if reducing the uncertainties requires a significant design change.

Of course the prime example of integration of physics and technology today is ITER [4, 5], and the ITER experience, past, present and future, is one of the foundations for the European DEMO design (e.g. see figure 4 of the EUROfusion roadmap, long version [6]). There is also much to learn from integration in new and recent facilities, for example JT-60SA [7], and others, e.g. [8–15] as well as from JET, the operating facility closest to ITER. There are also devices still in the design stage, e.g. [16, 17]. Finally, there is a range of other DEMO class power plant studies performed in various nations currently and in the past, for example in US (ARIES) [18, 19], Japan [20, 21], Korea [22], China [23]. While this paper is not intended to be a review, we note that, for example, the ARIES studies [19] address a wide range of integration issues and give the assumptions and rationale. They consider more advanced plasmas and/or technology than the European DEMO, which indicate ways to raise the overall efficiency, etc of the plant. The international studies as a whole contain a broad range of ideas, including digital approaches [22].

The scope of this paper is as follows. ‘Physics’ here refers primarily to plasma physics, not to the similarly rich and important field of materials physics (although we make some passing references to that). First, we give a set of major areas where physics and technology/engineering interact. Then we describe three examples in more detail (including some of the uncertainties): using advanced divertors to address power exhaust; effective ways to manage the large tritium throughput while keeping the overall tritium inventory low and finally elements of a controlled plasma scenario integrated with the plant systems. After that come some examples of integration in existing/imminent facilities followed by a discussion of uncertainty types and how they might be handled. Finally, we outline some integration tools and their forward path before the overall summary. This paper is aimed to inform generally and provoke thought, not to be a comprehensive review.

2. Major science and technology interfaces

Figure 1 is a reminder of the scope of the full integration in ‘space’ which is generally the starting point.

There are several critical areas where the physics of the plasma (and materials) interact strongly, directly or indirectly, with technology and engineering and vice versa. There are other critical areas which also have to be considered for a fusion plant. We populate these two groups below. In the first group we mention [in brackets] some consequent technology aspects/interfaces and impacts on the engineering design, and then expand a little. Three examples are explored in greater detail in section 3: exhaust, tritium management and the integrated plasma scenario.

(a) Strong interactions between plasma and technology/engineering

1. Power balance of a fusion plant and choice of total fusion power, options for changing the power output [plasma operating modes and scenarios, efficiency of heat conversion components; component replacement frequency and maintenance duration]
2. Overall plasma parameters [toroidal magnetic field, radius, aspect ratio, elongation, fusion power etc]
3. Plasma shape, symmetry and divertor configuration needed for sufficient plasma performance [poloidal field (PF) coil positions and access for robotic removal of blanket and divertor components, ferromagnetic steels: ports/variation give 3D error fields]
4. Exhaust of power and particles: schemes and PFCs [flat top phase of discharge, ramp-up and down, control of transients, erosion, material properties under irradiation, control and blanket armour impact on tritium breeding]
5. Fuelling, gases to assist power exhaust, pumping of helium [tritium inventory, release rate of T from blankets, time to separate species and isotopes, space for pumps, conductance]
6. Plasma control and the systems, their power and energy needed to manage transients [energy stores, demand on grid, need for internal coils to reduce power demand, ports for heating and current drive systems, pellet access through blankets, disruption control; integrated commissioning strategies]
7. Off-normal events especially disruptions or confinement collapse [proximity of scenarios to disruptive sequences, condition of the PFCs to minimise fragments, controls to avoid and mitigate, design to minimise consequences and allow fast recovery; additional structural components of the breeding blanket which can reduce the tritium breeding ratio; internal coils for control of divertor reattachment]
8. Variations in plasma performance/fusion power, both the steady levels and slow and fast transients [resilience of components and systems to variation in load conditions].




Figure 1. End-to-end spatial integration based on ideas for the European DEMO⁵. The photograph at the top left is of a prototype tungsten monoblock assembly (dimensions of a few cm) for a divertor target. Behind all these will lie an integrated supply chain of human and industrial capability of a scale and diversity seldom seen. The DEMO building schematic is reprinted from [24], Copyright (2018), with permission from Elsevier.

(b) Other areas critical for an integrated design:

1. Integration in time

- (i) over a pulse: start-up, burn phase, ramp down, pulse length, control
- (ii) over neutron dose: material and component performance change
- (iii) over design and build: coherent integrated concept
- (iv) over operational programme: plasma scenario & control development (commissioning may be radically different from today’s tokamaks); operational phases; decommissioning and waste management

2. Tritium self-sufficiency for given blanket coverage after divertors and penetrations [lithium enrichment, neutron multipliers, volume of structural material; neutronics calculations and their accuracy, tritium release and retention rates]:

3. Safety, environment and regulation [containment structures [25]; tritium boundaries, including penetrations and windows for RF microwaves and diagnostics; waste; management of internal and external off normal events]

4. Minimisation of activated and other waste/GWyr output [materials choice for activation; materials purity; remote maintenance in high γ -dose environments; materials/component lifetime—affects waste volume; P_{out}/P_{fus} , recycling and reuse of materials]

- 5. Architecture suitable for remote maintenance [number/size of blanket modules, number and size of cooling and other service pipes; topology, component mass]
- 6. Integration of the load assembly [containment structures, coils, vessel, blankets, divertor, cryostat etc [25]]
- 7. Integration into the buildings [space for services, pipes, coil feeds, enough waveguides etc; space for expansion tanks; paths for transporting components removed during maintenance etc [24]]
- 8. Cost of manufacture and construction [tolerances of components; complexity of components and systems; materials including their chemical purity and isotope mix]
- 9. Cost and schedule of supporting R&D, including dedicated facilities [level of experimental ‘demonstration’; scale of computational modelling]
- 10. Operating and decommissioning costs [component lifetime in steady and transient conditions]
- 11. Supply chains [materials, components, human capability].

Safety and environmental issues deserve special mention as they are fundamental to the device architecture, size of components, choice of materials, design of many systems and design of the buildings to allow safe robotic movement of large, activated components around the plant. The starting point is that during normal operation and maintenance there will be comprehensive shielding of neutrons, x-rays, gamma-rays; tritium containment and management of

⁵ DEMO figures from Top level Roadmap (euro-fusion.org) [6, 24], DEMO unveiled—Fusion for Energy (europa.eu).




Figure 2. Indicative (not calculated) power flow for a tokamak power plant. Some concepts consider using the low-grade heat to improve the overall efficiency, e.g. for pre heating of heat transfer fluids, and potentially thermomechanical not electrical energy for some of the recirculating power (e.g. some classes of pump). For EU DEMO, P_{fus} is presently set around 2GW aiming at several hundred MW electricity output, in pulsed operation [2].

activated material will be scrupulous; all maintenance in active areas and many other operations will always be conducted fully remotely. The building design will be affected by the likelihood and impact of external events (such as seismic activity and flooding). How this is converted into regulation depends on a number of factors including public opinion as well as more formal quantitative analysis. The regulatory environments adopted for fusion will have an important impact on the design and operation of the plant⁶.

In principle all of the above (and others) need to be considered simultaneously to avoid optimisation of one making the challenge for another insurmountable, and conversely to allow one challenge to be eased by generating margin in another area. Some examples can be discovered later in the paper. In practice resource limitations mean not all areas can be pursued in full simultaneously so choices may have to be made, with attention to the strength and impact of the interactions. However, it is unlikely any can be ignored, and the management will need to be dynamic: some interactions may weaken, others may become critical as the design and understanding develop.

Most of the integration issues are recognised and included in the EUROfusion DEMO programme, e.g. the Key Design Integration Issues, KDIIIs [26]. While there are other important national, international and private industry programmes, reference will often be made to the EUROfusion DEMO programme as an example [1, 2, 3]. However there are some

radically different approaches to integration, for example the ARC (affordable, robust, compact) concept [27].

2.1. Power balance of a fusion plant and choice of total fusion power, plasma operating mode

The ultimate goal of providing sufficient net power (heat or electricity) is an overall integration which affects the plasma solution via the total fusion power and the recirculating power needed to maintain and control the plasma. Although in principle the efficiency might not be a critical criterion if there is not a free market (e.g. due to political choices to ensure sustainable and reliable low carbon baseload power) there are several reasons to minimise the number of fusion power plants for a given total power, not least availability of some materials (due primarily to the supply chain and cost, but potentially also natural abundance). Looking at figure 2 it can be seen that recirculating power has a major impact—each MW of power injected to sustain or control the plasma could need ~ 5 MW of fusion power (heat to electricity conversion then the losses of converting the electricity back to injected power—higher efficiency plasma heating and current drive systems clearly help). The actual plasma Q ($=P_{\text{fus}}/P_{\text{injected}}$) needs to be much higher than this factor 5 to balance all the other losses and typically $Q = 30\text{--}40$ is targeted, which means $P_{\alpha} > 6P_{\text{injected}}$, i.e. the plasma is strongly self-heated and thus self-organising. A plant with high grade heat as the main output would have a different optimisation but still may well need to generate electricity for the on-site systems.

The choice of plasma scenario affects the overall power balance in several ways (and vice versa), depending on the

⁶ JET is managed under a non-nuclear regulatory regime, ITER a nuclear one. The nature of the fusion regulatory regime for power plants is not yet known, but active studies and consultations are underway.

arrangements with the external energy grid at the location of the plant:

- (a) The total fusion power to ensure sufficient net output, after subtracting the recirculating power
- (b) The power and energy needed to handle plasma transients—this could come from an energy store (and the energy in this store as the pulse progresses could imply pulse termination if there is not enough reserve to handle future transients).
- (c) Steady state current drive can take a large recirculating power but there is a virtuous circle, scenarios with higher bootstrap fraction need less current drive, and have higher β_p and thus potentially higher fusion power, but probably at the expense of profile control challenges (or at least uncertainty).
- (d) Inductively pulsed plasma concepts (as for the reference European DEMO) could underly plants with continuous power output. Some inductive plasmas appear to need only low auxiliary heating and the plasma is expected to be much simpler to control (see below). However, the situation depends on whether power for start-up and ramp-down is available from the grid; whether the customers/users need continuous output power; how much power is needed to control fatigue in power conversion systems and turbines. AC operation has also been considered to improve the average power, with design implications. This is a part of handling uncertainty in both plasma performance and stakeholder requirements: we return to the choice of pulse length briefly later.
- (e) Adjusting to match demand: normally the goal is steady state supply, but some load following may be needed in some circumstances—see section 2.7.
- (f) The on-site energy stores needed for the aspects mentioned above (filling them is a recirculating power).

The total fusion power and power density has integration consequences on the overall availability and possibly other aspects:

- (a) The in-vessel components are expected to be replaced after a certain level of neutron damage (dpa, displacements per atom), or possibly activation. A 14 MeV neutron power flux of around 1 MW m^{-2} results in about 10dpa/full power year⁷. With a blanket lifetime of say 50dpa (the divertor may be rather less), at this power density components would need to be replaced about every five full power years, so for an overall availability of 80% (with steady state operation, assuming no other failures), 1.25 years is available for the major intervention to change the blankets. However at the same dpa limit and 3 MW m^{-2} from a higher power or more compact device, the replacement would need to be done in five months—this intensifies the drive for improved materials and/or designs accommodating degrading materials, as well as fast remote maintenance.

- (b) The regulatory environment adopted for fusion may affect the waste strategy and thus the balance between frequent module replacement at modest activation (from transmutation) or infrequent at higher activation (and neutron damage). It is worth noting that the power density probably does not have much effect on the total activation for a given choice of materials—to first order the activation and transmutations will be set by the total number of neutrons, i.e. the number of GWyr of fusion energy, although the waste volume will depend on the recycling and waste processing strategy.
- (c) Greater thermodynamic efficiency is possible with higher blanket temperature (or more precisely, higher temperature of the heat transfer fluid), and this may need higher power density.

2.2. Overall plasma parameters

The traditional approach to integrated design has been to use simplified, low fidelity plasma and engineering models to find a ‘consistent’ design point by means of generally 0D systems codes (e.g. [28, 29])—some precision in the device parameters has been found essential to allow the engineering design to make any useful progress, and there have not been tractable high fidelity models of the plasma and plant separately and combined. The underlying assumption here is that high order effects will not significantly change the main parameters and thus the design point is both valid and consistent between the plasma and the engineering—this assumption has to be tested. This approach was used for ITER, and subsequent research has upheld it to a large extent [30, 31]. It has also been used for EU DEMO [1, 2] and other concepts, and here when the integration is considered in more detail for the more challenging power plant environment, higher order effects can have substantial low order impact and a revised approach appears necessary.

While systems codes now have some 1D aspects embedded, the plasma modelling is still rather basic, usually with 0D confinement scalings used to normalise the transport, and the time dimension is generally missing. A modest advance gives further illumination and guidance: a more advanced but still reduced (and thus fast) 1D model of the plasma can be used with a simplified calculation to get the net electrical power as a function of the major device parameters—major radius (R) and TF B_T [32]. Figure 3 illustrates qualitatively for given values of aspect ratio (A) and elongation (κ); together these four, B , A , R , κ could be said to be the major engineering parameters, largely defining the exoskeleton of the tokamak.

At the heart of all these simplified approaches lie one or more 0D energy confinement scalings from experiments. There are two important points here: first the data is taken in the parameter regime of today’s or past experiments, not the target power plant and second, despite best efforts, the databases typically include a range of plasma performances within each device, different experimentalists, and different sets of hidden (or unaccounted) variables such as wall condition, power deposition profile etc.

⁷ This is strictly only valid for a specific element, ⁵⁶Fe in this case.




Figure 3. Indicative (not calculated) net electrical power in B, R space, based on ITER-like confinement scaling (using general features from [32] but allowing for some improved performance with higher elongation, for example). The curves are simple estimates for the field at the geometric axis for a given maximum field at the outer edge of the inboard leg of the toroidal field coil and allowing a gap for a shield and blanket, for an aspect ratio 3.1 plasma.

As a result of the second the database projections may best indicate the behaviour of a population of tokamaks with various hidden variables, not a single, optimised, device—an important distinction potentially leading to conservative and even pessimistic designs, which must be recognised. This point is also visible in existing experiments—it is well known that confinement scaling with some parameters (e.g. β) are different within individual tokamaks from the multi-machine scaling. An illuminating example is provided by the move from a carbon to a metal wall on ASDEX Upgrade⁸ and then JET. On both devices there was high performance with a carbon first wall and divertor, providing amongst others the basis for the performance estimates of ITER. The first wall was then converted to metals, Be and W on JET to mimic ITER, and W only on ASDEX Upgrade, more relevant to the European DEMO. In both cases the performance was found to drop substantially (primarily due to changes in the pedestal from the enforced operating regime to avoid melting divertor components)—this was unpredicted and not addressed in scalings (which do not include edge atomic physics and neutral penetration effects). However, intense experimental and model-based optimisation over a number of years has recovered reliable high performance on AUG [33, 34], and removed the deficit and indeed exceeded the carbon-wall performance in JET [35]. As with many plasma optimisations, it is also found that the time sequence within the discharge is critical, i.e. integration in time. A OD scaling based on the full dataset covering this evolution in performance with experience is likely to give misleadingly pessimistic predictions. This reinforces the benefit/importance of using deeper and more fundamental understanding as a basis of

the design especially when extrapolating further from today's devices, e.g. because the plasma is more self-organising due to the internally determined fusion power dominating externally determined auxiliary heating, or because the parameter regime takes models into untested regions. It is worth noting, however, that the degree of extrapolation varies considerably between concepts—in particular the approach of SPARC [17], the precursor to ARC [27], is to design the core plasma to be in the middle of the existing confinement database, transferring much of the extrapolation to the technology and engineering advances needed to make a compact high field device (there will still be issues relating to self-organisation of the alpha heating, the effects of alpha particles on stability and transport, and the exhaust).

2.3. Plasma geometry

It has long been known that shaping the poloidal cross-section improves plasma performance in particular by allowing higher plasma current for a given edge safety factor (e.g. q_{95})—confinement time generally increases with plasma current in present experiments, approximately linearly in conventional aspect ratio devices ($R/a \sim 3$). Stronger shaping (i.e. higher order, triangularity and higher moments) improves the edge stability, e.g. allowing higher or possibly more benign pedestals. The shaping is provided by PF coils and this is usually amongst the first integration tasks—how to arrange the PF coils to allow enough shaping while leaving large enough ports to remove large internal components (notably blanket and divertor modules). Having more coils also usually increases the forces between coils, requiring additional structures which also limit space for ports. Finally, the coils produce field components perpendicular to the TF coils leading to out-of-plane forces, again needing additional structures (very challenging in this case due to the size of the forces and the need to keep the strain in the Nb_3Sn or other ceramic superconductors to a very low level). These aspects are also factors in the exhaust design, as described in section 3. There are efficient tools emerging for optimising the coil positions to generate a particular plasma equilibrium, e.g. NOVA/BLUEPRINT [36, 37] and tools from CREATE [38, 39].

2.4. Exhaust and PFCs

Plasma exhaust has long been recognised as a central concept-defining issue for tokamaks. The European DEMO situation is described in [40, 41], and the future strategy in [42]. The main challenge arises from the large thermal power deposited by the fusion alpha particles (and the auxiliary heating for sustainment, current drive and control) and the small surface area of the plasma footprint on the divertor targets for the observed cross-field transport when scaled empirically from today's experiments [43, 44]. Furthermore, the heat transfer through solid plasma facing materials to the heat sinks degrades with neutron irradiation [45, 46]. As a result the power during the flat top cannot be handled solely by magnetic flux expansion and angling of the target plates (both increase the wetted area), so a solution is sought from several mechanisms and options, with some resulting complexity of design and then

⁸ Axially Symmetric Divertor Experiment Upgrade also known as AUG.

control (see section 3). Finally the exhaust has to be managed during the plasma ramp-up and ramp-down as well as the flat top: these transient lower power phases can, surprisingly, be very challenging.

2.5. Fuelling, and gas to assist power exhaust

The large majority of the hydrogen (deuterium and tritium) fed into the tokamak is needed to sustain the plasma density to allow a high fusion power and to provide gas to the divertor to dissipate the exhaust energy and assist divertor detachment; only a small fraction (a few percent) is to replace the fuel burned. The large throughput of tritium, perhaps several kg h^{-1} , appears to present a major challenge, not least in the starting inventory, but also in the inventory in the tritium plant to enable processing at this rate. However the integration challenge has driven new thinking, transforming the situation (see section 4).

2.6. Plasma control including off-normal events and disruptions

Plasma control through the pulse determines the viability of the scenario, perhaps fundamentally, and the number of observers (diagnostics) and actuators (especially heating and current drive systems). Transient variations in fusion power also have to be handled by the plant—similar approaches can be adopted as for better or worse than expected flat top performance (section 2.7). There are several other classes of challenging transients and off-normal events, e.g.: confinement transitions in tokamaks (L-mode to H-mode and the reverse) which impose demands on the radial position control [47, 48]; radiative or other core plasma collapses in stellarators [9, 49] as well as tokamaks; divertor reattachment and of course disruptions. Disruptions are very well-known phenomena, and a major concern given their potential for damage, and they have received extensive attention on ITER [50]. On the other hand, very well prepared and well-designed plasma scenarios can be operated disruption-free even at high performance, if there are no external events such as fragments of material entering the plasma, uncompensated failure of heating and fuelling systems, or incorrect control action due to diagnostics problems. They need to be addressed as part of design integration, and there are many approaches such as: designing a scenario where internal events (such as confinement transitions) can be controlled well enough to avoid a cascade of events leading to disruption; identifying root causes such as particular plant or discharge programming errors [51]; real-time identification of imminent disruption as part of a comprehensive disruption avoidance strategy; developing disruption mitigation systems; designing components and the vessel to survive disruptions (e.g. by ‘sacrificial’ limiters [52]). On ITER the imposition of a particular (fission-like) regulatory regime has reached back into tokamak physics in the case of disruptions, due to the risk of damaging the vacuum vessel especially given the uncertainty in the behaviour of halo currents. The approach to this uncertainty is interesting: as well as enhancing the mechanical structure, there has been intense effort on finding robust

approaches to mitigate disruptions [50] so that the situation of highest uncertainty (sustained large, slowly rotating halo currents) simply does not materialise—in principle this may allow the desired scenarios to be used without compromise.

2.7. Variations in performance and fusion power

The uncertainty in plasma performance could result in a range of operating points—higher or lower than planned, and there can be transient variations due to disturbances in the control loops. There may also be a stakeholder request for some degree of load-following beyond that which can be handled by local energy stores, i.e. changing P_{fus} in a scenario, or allowing more than one scenario, so uncertainty management via multiple plasma scenarios (section 5) might overlap with meeting stakeholder requests for demand following. The main interfaces are:

- the first wall & divertor (i.e. the exhaust management and systems);
- the breeding blanket’s other two roles: converting the neutron power while keeping the materials (especially steels) in a good operating region, and shielding the vessel and other internal components (such as coils) for sufficiently long life (which may be set by calendar years as well as GWyr);
- the maintenance system and waste strategy (short lived in-vessel components);
- the balance of plant (heat exchangers, pumps and generators) where a single design may struggle to handle significant variations in time or operating point.

There can be various approaches for different cases. If the flat top performance is worse than expected this can ease the exhaust and increase component lifetime, but is likely to reduce the plant attractiveness, so the design should have margin to recover P_{fus} , such as capability for higher plasma current (subject to optimising the recirculating power) or alternate scenarios (see section 5). If the performance is better than expected (higher P_{fus}), it would be good to take advantage of this, and add engineering margin to accommodate. However if the balance of plant cannot handle it (and it is not possible or chosen to add additional heat exchangers and turbines) or the exhaust becomes unmanageable, then the power has to be reduced, e.g. by lowering the density, reducing the tritium fraction, or lowering the plasma current which may make the scenario easier to manage and allow longer pulses due to the larger non-inductive fraction possible (and/or lower recirculating power).

These alternative scenarios have to be feasible (e.g. exhaust at lower density). Similar considerations apply to handling transient excursions in fusion power. See section 5.

3. Integration example 1, exhaust: plasma, materials, technology and engineering interaction

For a tokamak at least there are two primary tasks for exhaust and six basic ideas to accomplish them (many of these also translate to the stellarators). Their integration brings several players to the field (figure 4): the core, pedestal and divertor




Figure 4. Outline of the integrated plasma exhaust situation. P_{SOL} is the power crossing the separatrix into the scrape-off layer.

plasmas, plasma facing materials and components, engineering and remote maintenance, the pellet and gas supply and the pumping systems (see also the section on tritium management below). To provide some background, the power to be exhausted in a 2 GW $Q = 30$ plasma is nearly 500 MW, compared with 150 MW for the ITER base case. For a 9 m DEMO with 70% of the power radiated to the first wall, the average first wall power density is $\sim 0.2 \text{ MW m}^{-2}$, but this will have substantial poloidal variations [40] and neutron attenuation by the necessary armour becomes significant at these levels. The corresponding divertor power is $\sim 140 \text{ MW}$ which still requires substantial attenuation and spreading to match credible PFCs. Should the radiated fraction drop to 60% (due to control limitations for example) an additional 47 MW must be handled by the divertor—this on its own is enough to substantially exceed the capacity of existing PFCs in a conventional divertor configuration without dissipation. One might ask why the radiated power cannot be taken closer to 100% to ease the divertor problem: there are two reasons why this is undesirable: first the point just made about the impact of small changes in f_{rad} but *a fortiori*, and secondly, experiments show that a significant power is needed to ensure a pedestal is created and sustained (this is the L–H power threshold for H-mode, but no-ELM⁹ pedestals such as QH and I-mode also have a power threshold in today’s experiments). Although we do not yet know the requirements to make a pedestal at DEMO parameters, it is likely to need a substantial power. Of course, if an improved core plasma could be found that does not need a pedestal the situation would change; there is no support for that yet for the transport calculated and measured in ITER-like core plasma scenarios, but there are some indications for negative triangularity plasmas.

⁹ ELM: Edge Localised Mode.

The two tasks are, simply expressed:

- to remove the helium produced by the fusion reactions and control impurity levels
- to handle the heat deposited in the plasma by the fusion alphas and the auxiliary heating and current drive systems (to sustain and control the plasma).

The six main ideas are

- create a divertor, with sufficient neutral density to channel the helium to the pumps
- radiate much of the power from the main plasma to the chamber first wall (using seed impurities) while minimising sputtering erosion by fast neutrals (hydrogen and higher mass impurities)
- identify plasma scenarios without ELMs or other significant unplanned slow and fast transients (yet with sufficient confinement)
- dissipate the power transported to the SOL (scrape-off layer) before it reaches the divertor target PFCs (this usually involves a detached divertor plasma)
- expand the plasma channel to spread the power to the divertor targets; configurations with more divertor legs
- develop materials and components for the main chamber and divertor PFCs to handle higher power even after neutron irradiation.

Integration considerations include:

- consistency of the divertor plasma with main-plasma scenarios/configurations that avoid ELMs, including transients due to fuelling pellets and changes in the auxiliary heating due to systems switching on and off

- management of power between the different divertor legs and targets, and to the divertor throat (e.g. for x-point radiators [53])
- generation and control of dissipation in the divertor (i.e. detachment) with hydrogen¹⁰ and impurity gases, and if necessary internal coils to allow fast sweeping of the strike-points in event of reattachment
- exhaust power handling in the ramp-up and ramp down phases, whether for limited or diverted plasmas (and thus the relative timing and control of the fast ramp of alpha heating [54] and plasma detachment)
- particle transport in the core and pedestal that allows hydrogen to reach the centre while helium and other impurities leave (relates to the turbulence type and regime in the plasma scenario)
- location and capacity of the pumps in the divertor(s)—there may be more options for the pump aperture location with detached plasmas
- consistency with the throughput capability of the tritium cycle (see section 4)
- ensuring the first wall armour is thin enough not to degrade the tritium breeding too much
- matching the required heat extraction from the PFCs with properties of coolant fluids and materials properties after irradiation at the coolant temperature
- location of PF coils to make the divertor configuration while allowing rapid remote replacement of the blanket and divertor and avoiding excessive forces and structures or reduction of the solenoid flux swing
- sufficient tritium breeding and neutron shielding of the TF coils, given the divertor replaces some of the blanket and will need thick shielding behind it
- identifying and accommodating diagnostics and control observers (model-assisted signals are hoped to reduce diagnostic requirements compared with today's experiments).

This list is quite long, but actually merely describes the present integration activities (not a new set of challenges). The elements are all known and most are being addressed actively, e.g. [41, 42, 52]. However, whether there is a viable DEMO solution with an ITER-like single null divertor is still uncertain and will remain so for some time: this uncertainty has triggered intensive efforts in several programmes, notably EUROfusion's where an organised exploration of alternative exhaust concepts was initiated, as a back-up. This used a range of ideas for advanced configurations originating from around the world, the so-called snowflake, super-X and X-divertors to name the main ones, as well as double null configurations to try to spread the power amongst more divertor legs (also a feature of the snowflake)—some outcomes are described in [55] and with greater emphasis on engineering integration in [56, 57]. In addition, a programme of work on advanced PFCs and materials was launched, notably liquid metal concepts, alongside improvement in the ITER tungsten monoblock approach.

This exercise is an interesting example of taking imaginative plasma physics, engineering and materials ideas and exploring how they evolve as the integration realities are imposed.

The first alternative divertor configurations had large numbers of PF coils which had two problematic consequences: (a) large forces between coils, as the currents were large and with large coil-coil variations to achieve the higher order multipoles and large low field regions needed for some configurations and (b) constraints on access—the space between the coils did not allow large enough ports for removal of entire blanket segments (the present reference scenario for the European DEMO)—this was compounded by the mechanical structures needed to withstand the forces. Constrained optimisation approaches [37, 38] allowed progress with fewer PF coils and thus access for remote maintenance [56] (whether this access is sufficient is not yet confirmed [58]). The next challenge was in the TF coil: the larger divertor regions require larger coils¹¹. This then pushes the PF coils further from the plasma, leading to increased currents. To allow closer PF coils, changes to the TF coil contour were explored, not the traditional bending-free D shape (which has anyway been modified for the reference EU DEMO coils)—this adds complexity to the TF coil design and manufacture to keep the strains very low (as required for ceramic superconductors such as Nb₃Sn). Finally, the new configurations have larger fields perpendicular to the TF coils—these lead to larger out-of-plane forces and mechanical structures to resist these (metallurgical variations of very thick forgings need managing [59]). The overall magnet design needs further work (large high precision magnets are always challenging especially with large forces), although there is optimism that there is an acceptable solution. See figure 5 (from [56]). It is immediately apparent that engineering integration has taken the 'super-X' configuration far from the original super-X concept [60]. Finally, while the reduction in tritium breeding areas appears manageable, they also need additional shielding to restrict neutron heating of the TF coils. There is thus a tension between the pressure from the engineering axis to make the divertor as small as possible and the goal of providing a resilient exhaust solution.

In parallel with all this one has to see if the now constrained configuration actually helps the exhaust challenge enough. This is addressed for the 'super-X' configuration in [56, 57, 61] where it is seen that the additional divertor volume and connection length is predicted to allow much more power to be dissipated if the neutral density is high enough (see the tritium integration section below). An additional advantage of long divertor legs, especially with variation of total magnetic field along the leg (total flux expansion, most easily achieved in a spherical tokamak (ST)) is an increase in the 'detachment window' (e.g. range of power for which the divertor remains detached and the detachment front does not enter the main

¹⁰ In practice this may be the same deuterium and tritium mix as for the main plasma—see section 4.

¹¹ Several tokamaks have had two- or more part TF coils allowing internal PF coils, e.g. C-Mod, COMPASS-C/D, MAST/MAST-U, NSTX/NSTX-U, START. This is more complex for superconducting coils but is being explored for ARC and STEP in particular.




Figure 5. Cross sections for some of the alternative divertor configurations considered for the European DEMO showing how the PF coils and ports and been co-optimised, and how the TF coils deviate from the bending-free D-shape even for the reference SND. Reprinted from [56], Copyright (2021), with permission from Elsevier. TF coils and inter-coil support structures and divertor cassette in grey, PF coils in blue, breeding blanket in purple, vacuum vessel and ports in green (see the on-line version for coloured plots). The scales are different, the plasma volume is the same in all cases (2350 m^3). The region around the X-point is sometimes called the divertor throat.

plasma too strongly, if that is a problem, see below) [62, 63]—whether this would be enough to avoid the need for sweep coils to mitigate the effects of loss of detachment is not yet known.

Originally the super-X was assumed to be in double null to reduce the power to the more difficult inner leg, but the DEMO studies suggest that the inner leg can be detached even in single null, so the complexity of managing detachment in four divertor legs might be avoided. The conventional DN and the snowflake do have four legs to share the load, bringing a requirement for both magnetic and detachment control for the extra legs [56, 58]).

The double null and snowflake configurations are more complex to model (since they have two or more active X-points), so the modelling data is more limited so far.

Viable control is an existential issue for any exhaust concept. There are three main elements:

- Control of the power entering the divertor and incident on the first wall: measurement and management of the radiated fraction in the main plasma is not straightforward (e.g. bolometer views are limited so overall precision may be low), and control of the core plasma may need variations in the auxiliary heating power (e.g. for control of neoclassical tearing modes) so a divertor that is resilient to variations could be essential.
- Control of the magnetic configuration: this is especially important for configuration with more than two divertor legs: e.g. snowflake and double null (conventional or SX), where the power to the different legs can change greatly with small changes to the plasma vertical position for the DN (except perhaps for a double x-point radiator scenario), and changes in internal parameters such as β_p , l_i for the snowflake. The strike-point positions of the single null SX are also sensitive to changes in β_p , l_i .

- Control of detachment: the energy sink of a volume of hydrogen and impurity gas/plasma is modest and even small perturbations can cause ‘burn-through’ or reattachment, although long leg divertors are more resilient [62, 63]. Controlling detachment in DN configurations is likely to be complex, not least due to intrinsic up-down plasma asymmetries (e.g. due to drifts) but also due to possible indirect interactions between the upper and lower detachment zones.

An important behaviour found experimentally in the reference SN configuration at aspect ratio around 3 is that the detachment zone readily moves to the X-point, i.e. most of the dissipation occurs close to or inside the last closed flux surface poloidally localised to the X-point and this has been shown to be controllable in ASDEX Upgrade [53].

There are several integration consequences of detached plasmas to consider:

- the main power load is on the PFCs near the X-point in the case of X-point radiation, i.e. the divertor throat (although the local power density on PFCs is presumably substantially less than for an attached divertor)
- the pumping concept is likely to be very different from a partly or fully attached divertor where the particles are guided magnetically to the pump throat (as on ITER)
- the neutral pressure in the main chamber may increase, raising the charge exchange neutral power loads on the first wall if the detachment front emerges from the divertor throat
- the edge plasma may be rather different for an X-point radiator configuration especially, the outer region of the pedestal will be two dimensional with strong poloidal gradients near the radiating zone—this may or may not help with a no-ELM pedestal of sufficient height

- potentially it may partly resolve the control problem in DN as the difficult-to-achieve precise magnetic balance of the upper and lower legs may be less important (even if this removes/weakens one option to adjust the up-down power sharing)
- the mechanical structure of the divertor (the cassette) might be smaller, with benefits to the breeding ratio, even for DN (but note that the divertor must manage ramp-up and ramp-down when the divertor may be attached requiring large divertor targets).

This all shows that the overall integrated optimisation goals, and PFC design targets, may be rather different from those considered in present-day devices, or might be chosen to be.

Interestingly at low aspect ratio (i.e. the ST), long outer leg divertors have an important extra feature due to the gradient in $\text{mod}(B)$ along the divertor leg, which should assist control of detachment and enlarge the detachment window [62]. Long leg divertors may also allow a different scenario optimisation, with lower upstream density (still with detachment) and thus more options for steady state with more efficient current drive. However, it is more likely that a double null (or close to double null) is needed in an ST to avoid overloading the inner legs (which are at small major radius so have low wetted area), with the associated control aspects.

An overall outcome of the exhaust integration studies so far is that there may be an attractive approach for an $A \sim 3$ tokamak that is in fact a relatively modest adaptation of the reference single null but with significantly longer divertor legs and increased heat exhaust capability [58]. Configurations closer to the original super-X appear to be attractive options for ST power plants. It also illustrates that fusion can indeed advance via innovative ideas adapted to meet integration constraints.

3.1. Uncertainties

The engineering and technology uncertainties can probably be addressed by existing methods, assisted by tools for the magnet optimisation. The plasma uncertainties are more significant, in the value of P_{SOL} , in the detachment control and detachment power window, in the compatibility with the main plasma scenario (e.g. a no-ELM pedestal, and an intense x -point radiating zone), and whether the increased neutral density to deliver the gains are compatible with tritium handling (see next section). There are experiment and theory programmes to address the issues, for example based around MAST-U (Mega Amp Spherical Tokamak Upgrade) [10, 11] which was designed to generate a wide range of configurations and divertor leg lengths, including a full super-X configuration (itself an integration challenge) and the advanced divertor programmes on TCV (Tokamak à Configuration Variable) [64], ASDEX Upgrade [65], DIII-D [66] and other devices worldwide. Providing confidence in time for initial decisions on either the reference single null or an alternative (longer leg explored here) will require ingenuity, or ways to allow significant design changes to be made late without adding too much

delay (that would need rapid and effective design tools). See also sections 7 and 8.

4. Integration example 2: how to provide tritium to the plasma while achieving an acceptable site inventory

Tritium release is the main radiological hazard of fusion plants during operation, and while T release may be a low risk for a well-designed plant, nevertheless the site inventory of tritium must be kept as low as reasonably practicable (ALARP), especially the mobilisable fraction. In addition, tritium is a globally scarce resource, so the required start-up inventory also has to be kept low.

Hydrogen (protium, deuterium and tritium) is needed for four main purposes: to sustain high enough fuel density for the required fusion power to be achieved; to provide dissipation in the divertor (and sufficient compression and entrainment of helium for continuous exhaust of the ash) while maintaining low impurity concentration in the main plasma; to replace the fuel that is burned (only a few% is consumed per pass through the tokamak) and, finally, protium (and perhaps deuterium) for the (re)commissioning phases at the start and after maintenance or other interventions. The requirements for dissipation in the divertor may in fact dominate the throughput, potentially requiring flows of several kg hr^{-1} , depending on the balance of puffing and pumping needed to maintain the required neutral density [56]. The main plasma is expected to be fuelled by pellets injected from the high field side (to take advantage of the outward ExB drift of the resulting plasmoids taking the fuel towards the core [67])—this needs pipes to take the frozen pellets through the blankets [68]. The D:T mix is assumed to be 50:50, although leaner burn may be possible reducing the tritium required for a given density, e.g. if there are sufficient fast ions to merit a revised overall optimisation including the tritium cycle. It is likely (see below) that most of the unburnt D-T fuel can be directly returned to the plasma. Fresh tritium is therefore needed to replace tritium that

- is burned
- is absorbed into materials or permeates into coolants (both can be recovered, but usually much later)
- decays (this can be significant in event of long maintenance periods or long retention in materials).

Short term sinks due to absorption into material and permeation can be a significant issue [69] (which also looks at global tritium management over the plant lifetime), and this is a key contributor to the initial inventory of tritium that needs to be imported to the plant. Estimating this start-up quantity accurately enough might need high fidelity in the modelling as well as significant detail in the design. The behaviour of tritium (and hydrogen generally) in materials is a complex scientific topic—and the take-up is affected by neutron damage making vacancies where T can be trapped. Recent theoretical and experimental studies for tungsten suggest the hydrogen trapping may saturate at low dose [70–73] but at a relatively large level (possibly reduced by adding small amounts of alloying




Figure 6. Schematic of gas and fuel injection and pumping/purification cycle (after [75]). Presently the separators are proposed to be in the divertor ports, before two sets of pumps (for the permeate and retentate).

metals such as Re [74]). In principle this might be delayed by implanting deuterium initially, filling some of the traps. There may anyway be some differential absorption and release of D or T.

The traditional image is of T and D supplied separately and combined into the pellets and gas feed, perhaps with the balance changed in real time as a part of the burn control. This works well for pulsed machines such as JET and probably ITER but is considered unfeasible for a continuously operating facility (whether cyclic or steady state plasmas) because the isotope separation is a slow process and thus requires very large plant (and associated high T inventory and cost).

This new integration issue is being addressed for the European DEMO by an innovative approach where the large majority of the tritium and deuterium is directly recycled from the exhaust back into the plasma bypassing the main tritium plant [75]. Assuming all the material and pump tritium traps are filled, the tritium plant now mainly has to provide the fresh T required and rebalance the injected fuel mix (e.g. if there is preferential absorption or permeation of D or T in materials). Buffers are still needed to deal with the increase in fuel and divertor injection rate as the density is increased and the divertor is detached and taken to high neutral density as the plasma is ramped up. The size of these buffers depends on the delay time within the direct loop, between exhaust from the divertor to ice in the pellet injectors and gas in the high pressure reservoirs behind the gas-valves. A possible drawback of the direct loop is that it means that the divertor gas has to be the same D:T mix as the main plasma, and thus the gaseous tritium at any time is increased compared with a deuterium-seeded divertor, but this is argued to be balanced by the reduced inventory and management of tritium in the main plant. The scheme is illustrated schematically in figure 6 (for more complete description see [75] and refs therein). For this approach

to work, technology developments are needed, in particular the system for fast separation of the hydrogen (D + T) from the rest of the exhaust gas in real-time. Presently the research in Europe is focused on plasma-assisted permeation through metal foil pumps (MFPs) [76, 77] that could in principle be placed in the divertor ports if the plasma-assist can operate well there given the tesla-level fields, and if a large enough area of metal foil can be achieved within the available port cavities. This approach places some demands on the main tritium plant to process the hydrogen that is not separated initially (the present target is 80% of the D + T being extracted from the exhaust gas and fed back into the plasma), but 20% entering the main T plant is considered manageable in steady state and it is far less than the demands without this direct recirculation. The D + T extraction fraction (for a given throughput) depends on the length/aspect ratio of the MFP tubes so is in principle adjustable, depending on the space available (the situation would be eased if the pumps could be outside the cryostats and/or bioshield, subject to conductance constraints). If one could get this from 80% to 100% (or very close), then the tritium plant can be sized essentially independently of the (quite uncertain) recirculating throughput. The tritium plant scale would then be determined by the fusion power and the (nominally) transient sinks in materials and the chronic permeation into coolants. Whether this 100% recirculation target is the best optimisation will depend on the engineering and technology needed, which of course depends to some extent on the throughput needed to sustain the plasma and exhaust. For the European DEMO a back-up option is retained: a set of multi-stage cryopumps operated cyclically (with the regeneration phases providing the required flow, the kg hr^{-1} mentioned above) which can also separate most impurity gases and helium from hydrogen, i.e. retaining the high recirculation fraction.

Since the tritium system is intimately linked with the whole gas and pumping system as well as the management of tritium in materials and coolants, there are other integration aspects including:

- efficient and rapid extraction of tritium from the breeding materials (solid or liquid)
- tritium required for start-up of the plant (and its minimisation)
- tritium needed to start other fusion plants
- handling the gas pulses from disruption mitigations systems (shattered pellet injection; massive gas injection)
- recovering tritium from coolant and heat transfer fluids (with reasonable power and space)
- effectiveness of anti-permeation barriers
- recovering tritium from components removed during maintenance and decommissioning
- recovering tritium from components *in-situ*
- tritium compatible pumps (and any side issues involved, e.g. if mercury pumps chosen)
- power and services to the tritium separators if in-vessel.

This short description shows how full integration reveals aspects not very important on today's experiments and ITER but critical for power plants, and has thereby driven innovation that can lead to a major simplification and size reduction of the large, complex and costly tritium plant. It is seen that the integration has to be studied in some depth and breadth to find viable and attractive solutions for power-plant-class facilities.

4.1. Uncertainties

The uncertainties in the fuel management include: overall throughput needed (related to the fuelling and exhaust situation); tritium retention and release from radiation-damaged tungsten; permeation and retention more widely; effectiveness and speed of tritium release from materials and fluids; development of reliable performance and capability of the metal foil pumps; the start-up tritium inventory and relation to the external supply. Timely quantification/resolution of these uncertainties is a factor if margin in pumping space and tritium plant capacity and external tritium supply are not adequate to handle all the likely situations.

5. Integration example 3: integrating an uncertain plasma scenario into an engineering design

Perhaps the central integration task is to find a full plasma scenario (end-to-end) that is consistent with a realisable engineering and technology design (or vice versa). As indicated in section 2 above, the most visible engineering choices for the core plasma are B , A , R , κ , i.e. the general features of the exoskeleton around the plasma (plasma current, I_p is not included, but the solenoid required for the selected operational mode will affect the major radius, especially for STs). However, there are many more aspects to the space-time integration and some of these will be explored briefly here—we have touched

on exhaust and tritium already. Some, including their uncertainties, have substantial impacts on the design and the overall programme. As already mentioned, it can be misleading to suppose that an approach starting with the 0D parameters via scalings will lead to a consistent solution—evidence to the contrary has already been found in the European DEMO programme. A solution cannot be adequately described by even a full set of 0D dimensionless parameters (including β , q_{95} , and triangularity δ) as well as dimensional parameters, and adding a few profile-related parameters (e.g. $q(0)$, q_{\min} , l_i) is only a modest improvement. Furthermore, these do not address two existential aspects: controllability of the flat top scenario and controlled access (and termination), and these should be considered from the outset. Finally, the uncertainty in the performance (and control) must be addressed. These can all have substantial implications and could eliminate some concepts, or conversely drive designs that allow for more than one plasma scenario. Figure 7 illustrates very generically the situation for a concept that has three potential scenarios with different characteristics.

Essentially all tokamak plasmas considered for power plants are vertically unstable—this is a very well-known issue and routinely managed in many tokamaks, the instability growth rate is set primarily by the plasma elongation and the general shape of the current profile (parameterised by the internal inductance l_i) together with the location, inductance and resistance of the metal structures surrounding the plasma. The control is by coils and their power supplies. There is an integration issue arising from the power and response time in the feedback system, and ITER has been forced to have non-superconducting coils inside the vacuum vessel; this may also be needed for fusion power plants, given the high instantaneous reactive power needed if the coils are external [78]. While adding complexity, this is a manageable and readily characterised approach (coils can be very well-shielded on the present European DEMO and thus probably last for the lifetime of the device).

Starting with the 'flat top' (what has to be held constant and with what accuracy is a discussion topic), examining control can be an effective way of assessing and managing the inherent non-linearity and internal coupling of a plasma, and also provides a mechanism to handle some classes of uncertainty (section 7). For this to work, the underlying model has to be sufficient, with enough of the interdependencies—e.g. pedestal stability depends on details of the plasma shape, as well as the state of the core plasma [79]. There are apparently enticing advanced steady state scenarios but where the transport depends on the current profile and there is a large bootstrap current whose profile depends on the pressure gradient which in turn depends on the transport. In an ideal situation this self-organising situation would have natural feedback so it is thermally stable around a high fusion power point, but it is not obvious that will happen (some advanced scenarios have such good transport that the pressure profile steepens too much and drives large scale instabilities, and the pulse can terminate with a disruption). How/if these control schemes are to be developed and commissioned in plasmas without tritium is another integration question—flight simulators will be




Figure 7. Indicative/hypothetical landscape of plasma scenarios in a particular design. It may be possible/desirable to allow for more than one scenario in a given concept. This diagram represents a hypothetical 2D projection of a multi-dimensional space for discussion purposes—a diagram for a real device may be very different.




Figure 8. Schematic of the environs of an operating point. Three general types of plasma behaviour are illustrated—one is self-limiting (e.g. critical gradient transport) and one is ‘self-improving’ (and potentially thermally unstable) but with higher performance, and a third has a very narrow operating region which could be harder to maintain or reach, but with higher fusion power. Each will have an acceptable range of fusion and exhaust power set by the wider goals and system capabilities, and each will have a control range before the scenario is lost, due to plasma or thermal instability, disturbances from system failures, some irreversible change in transport or impurity content, etc. It may be possible to access more than one in a design (deliberately or by chance). Note that although superposed on same scales here, the scenarios may be well separated (e.g. as in figure 7). Note also that this is an illustrative representation of what will in practice be a multi-dimensional space.

key, see below. Figure 8 illustrates generically some different potential types of plasma scenario with different control requirements. Actuators also need to be able to compete effectively with the strong alpha heating.

Fortunately, there are scenarios which are much more manageable: they may have lower power density than some concepts, but recall the comments above about the link between higher power density and shorter materials life. The European DEMO presently considers a pulsed design with a low bootstrap fraction and a core plasma regime with ‘stiff’

self-limiting core transport (like the ITER baseline scenario), i.e. where there is a critical gradient above which turbulence becomes much stronger. This means that if plasma is chosen such that the self-limiting profiles are stable to major instabilities (or they can be controlled, e.g. neoclassical tearing modes via localised electron cyclotron current drive) and the fusion power is large enough, the control is relatively straightforward, and the main concerns are in the control systems themselves—e.g. occasional failures of the pellet injectors or individual gyrotrons in a system with several tens of gyrotrons. This has implications on the reliability required of each system and the level of redundancy. Redundancy is however a rather cumbersome way of handling poor reliability, in terms of cost, complexity, additional failure modes as well as space around the tokamak, although some redundancy can be achieved via dual function systems. System reliability is therefore an important part of an integrated design and is quite likely to be quantifiable and (in principle) tractable, especially for the more independent components such as gyrotrons.

Even in this attractive scenario the pedestal remains a challenge—the standard H-mode pedestal is an example of a self-organising region that drives towards instabilities (ELMs) considered unmanageable for the pre-conceptual reference European DEMO [80]. There are several ideas to address this, but all have uncertainties either in their likelihood of success or their impact on the overall performance for a given core scenario and device size. If the device size has to be constrained, or more advanced and complete calculations show that the core transport limits the fusion power below the overall requirement (for an ELM-less pedestal), then alternative scenarios for the same device may be needed as indicated generically in figure 7—that is anyway a desirable feature if both pulsed and steady-state need to be considered [81].

There is one more integration aspect from the core plasma performance: even for the relatively conservative self-limiting scenario of the European DEMO. If, for control reasons, the plasma must be operated at a ‘sweet spot’ with fusion power set by critical gradients, uncertainty in transport and the value of the critical gradient could lead to higher as well as lower




Figure 9. Schematic of the time evolution of a tokamak discharge and some of the factors to consider when integrating into a full fusion plant.

power relative to the reference prediction. The rest of the plant must accommodate this band, as well as temporal variations outlined earlier. As indicated above this can perhaps be handled for a fixed balance of plant (heat exchangers, generators) by designing to allow a range of plasma currents (and probably pulse lengths, for a pulsed device).

Moving next to the whole pulse trajectory, figure 9 illustrates some of the factors that need to be taken into account. For example, considering the PFCs [52], when does the divertor need to be detached? Controlling detachment is challenging when the power crossing the separatrix is increasing rapidly as the fusion power ramps up (over around 10 s for ITER simulations [54]). Would it help to have the strike-points for the attached phases on a different type of PFC with a larger inertial capacity? Assuming a pedestal needs to be created, should this be done before or after the plasma is detached and does that change the physics of pedestal formation? The ramp-up and ramp down will need substantial power for the PF coils and auxiliary heating systems, perhaps totalling well over 100 MW for a significant time (e.g. many minutes). This has to be provided either by the grid or from large local energy stores (which need to be filled, a recirculating power).

Control brings other integration elements into play, for example:

- sufficient ports for all the systems (many tens of high power microwave waveguides for gyrotrons)
- integrating ICRH (Ion Cyclotron Resonance Heating) antennae into the blanket if ICRH is included, and port-based systems are not adequate
- physical space and extension of the tritium boundary for neutral beam systems if included
- the electrical power required for the systems, often high pulsed demand—e.g. thermal or magnetic energy stores—which have to be filled and maintained to avoid a limit to the pulse length of the tokamak if the store is exhausted

- measurements systems (diagnostics) compatible with the environment and fast enough to act as sensors for control [78] and models to create the ‘observers’ for the controllers.

The EUROfusion approach to developing a consistent scenario is outlined in [80, 82]

5.1. Commissioning of a fusion plasma—a different kind of integration

A final aspect to consider is the **commissioning** phase of the power plant. This may be very different from today’s tokamaks and indeed ITER. Consider the final plasma—for $P_{fus} \sim 3\text{GW}$ and $Q = 40$, 600 MW of alpha heating will be supported by 75 MW of auxiliary power. How does one prepare sufficiently for this? Commissioning the individual systems is presumably not enough, likewise commissioning them together with no plasma, or a low power one. On the other hand, if one were to attempt to replicate the fusion plasma in full in protium (to avoid activation) external heating of 675 MW would be needed, requiring electrical input of perhaps 1.3 GW (for around 50% efficient heating systems), or more if confinement in protium is worse than in DT (as expected from today’s experience). In addition, advanced control would be needed to simulate the parameter-dependent alpha heating profile. It is often assumed replacing the alpha heating with auxiliary heating is impractical for various reasons—the cost and disturbance will indeed be large, but implementation could be feasible if the option were included in the design from the outset and it will probably be a question of risk appetite (see below). This is part of a range of considerations, e.g.:

- time spent (deferring the net output from the investment)
- activation of the tokamak (including from deuterium operation, if needed) at a time when modifications and repairs

- may be needed to address discoveries or early component failures
- uncertainties in the plasma scenario (e.g. to ensure no ELMs) due to the new plasma regime or changes in the boundary conditions which may need extensive runtime, and special tools and diagnostics
 - diagnostics for the DT phase may not give adequate (or any) information in a protium commissioning phase (e.g. neutron and gamma diagnostics) and conversely diagnostics needed to help commissioning may not survive significantly into the DT phase
 - how/whether to simulate experimentally the intrinsic nonlinearities due to the alpha heating
 - how to attain high beta if strong bootstrap current is required in the final DT scenario
 - how to develop control at reduced plasma parameters in a way that gives confidence
 - how to develop disruption avoidance and management convincingly at reduced parameters (to reduce potential damage to the device during development).

For a concept where the plasma extrapolation is modest and the confidence in the plasma simulators is high the commissioning uncertainty is lower, although there are generally behaviours that depend on the device as actually built. In particular, plasma performance control and disruption avoidance schemes (for tokamaks) might preferably be developed before tritium is introduced. Where the concept is more enticing (e.g. higher efficiency) but the extrapolation step is larger and/or the confidence in the plasma simulators is lower, and where the role of alpha heating is likely to be stronger (e.g. driving large bootstrap current fractions), the stakeholders (including the operator) may wish to invest funds and time in high performance non-active operation (i.e. a one-off addition to the funds and energy consumed in construction). Other options include gradual introduction of tritium; this would mean any plant modifications would need full remote handling from very early on, as well as assessment of any nonlinearities in the tritium cycle, and also that a full breeding blanket would be needed (conceivably the breeding blanket, or the breeding material, could be installed after a protium or deuterium phase).

ITER adopts an approach of using hydrogen (protium) and helium plasmas before deuterium and then DT operation, over a period of several years. How or whether this is transferred to power plant commissioning is a key question which probably needs to be addressed early on (cost and schedule impact vs risk mitigation).

To summarise, options/ingredients include:

- protium-only commissioning to avoid activation, perhaps before the costly breeder blanket/breeding material is installed
- enhanced auxiliary heating to simulate alpha heating and generate bootstrap current (the systems could provide spares/redundancy for the high P_{fus} phase, e.g. gyrotrons and power supplies) as a one-off addition to the cost and energy of construction and commissioning
- operation at lower current and TF to allow dimensionless parameters such as β_N , q to be attained at feasible auxiliary power (low B_T may need changes to heating schemes such as ECRH)
- additional diagnostics to provide greater information, and calibrate the DT-phase diagnostics (noting that several rely on fusion products such as neutrons and gammas) and the models used to compensate the sparse data
- phased introduction of tritium, noting that the device will activate rapidly, and a full breeding blanket will be needed
- early transition to full DT operation with heavy reliance on flight simulators (see below)—the protium phase might be mainly for system shakedown
- use of flight simulators to bridge the gap from H to DT.

Some of the choices may be guided by new regulatory arrangements being considered (potentially different from those on ITER), and by a potentially higher risk appetite of some stakeholders (e.g. more reliance on modelling and flight simulators than would be the norm today).

5.2. Uncertainties

several are indicated above, arising from the differences in the plasma parameters from today's experiments, gaps in the modelling and understanding (epistemic uncertainties) to bridge the parameter differences and the speed with which high fidelity models (and their reduced fidelity-conserving versions) can be developed. There is uncertainty in the willingness of stakeholders to accept the risk of a large behaviour gap between the commissioning and final plasma, or the expense of commissioning with high auxiliary power. Possible strategies are progressively higher fidelity models and flight simulators (see sections 7 and 8) and concepts that can handle a range of performance levels and more than one scenario (e.g. [81] indicates elements of how plasmas with different currents and pulse lengths, from short pulse to steady state, might be accommodated in a single device at the design stage). See also section 7.

6. Integration in today's facilities: converting physics ideas into operating hardware

Every existing tokamak and stellarator is an example of integration of science, technology and engineering, the most significant today being ITER, with JET the most relevant operating plant since it combines remote maintenance and tritium with a high performance tokamak. ITER breaks new ground in the scale and complexity of integration, and in particular relies on new supply chains of industrial capability and advanced techniques, and of materials, notably Nb_3Sn superconductor wire. Building adequate supply chains of materials, industrial capacity and skilled STEM individuals will be a critical factor for power plants.

The core ITER concept is based on extensive experimental databases, combined with significant but, compared with some concepts, relatively modest extrapolation from existing experiments and regimes. Other devices have a stronger reliance on theory. A more adventurous example is Wendelstein 7-X

where the complex modular coil set is based on a theoretical optimisation of particle orbits to reduce neoclassical transport (potentially very large in stellarators), and it was not initially obvious that the twisted superconducting coils could be manufactured with sufficient accuracy and supported adequately; indeed they proved challenging, but the device was constructed successfully and the results are very encouraging [8, 9, 83]. There are other important examples most notably JT-60SA, and another case of a theory-based alternative design is the MAST Upgrade ST, where the Super-X divertor concept was converted into reality in the design and then manufacture, construction and first results [10–12], albeit again with challenges on the way.

These integration examples have shown the complexity but achievability of integration without significant compromises to the physics design and provide an important platform from which to address the more conflicted integration of power plants with their additional systems and constraints. Their approach to uncertainty is however fundamentally different since they are research/exploratory experiments, whose *raison d'être* often requires significant scientific uncertainty and hence flexibility—recognising this difference is an important step.

7. Uncertainty estimation and handling

7.1. Why uncertainties matter at the design stage

There need to be ‘sufficiently’ reliable predictions of the performance of the whole plant in order to allow major investment and other decisions to be made. While the decisions will be influenced by the risk and uncertainty appetite of the investors (which, as mentioned above, may in some cases be much higher than that of the research community), some estimate of the performance uncertainty is surely required, for internal as well as external stakeholders.

There are uncertainties in the plasma, materials and technology performance and behaviours, which affect, perhaps fundamentally, the integrated solution, e.g. determining the major engineering parameters (B, A, R, κ), hence size. Handling these is one of the greatest challenges of designing a fusion plant, especially when technical and cost constraints limit options for design margin and flexibility. On the other hand, uncertainty can also create opportunities—almost by definition it means that performance of a part can be higher as well as lower than expected. Time is another dimension: the top-down schedules (from stakeholders and externalities such as climate change) mean that design choices have to be made in the face of uncertainty and this is especially testing in the case of gaps in extrapolable understanding as well as experimental data.

The predictions and their uncertainties both require models, and today the models are incomplete in various ways. How to choose between incomplete models is an interesting topic, since they will be used in new regimes, where better or worse fits to today’s experiments data may not be adequate criteria, especially if the model has free parameters (this relates to the deep topic of the meaning of ‘validation’).

This whole area has had relatively little organised attention in the fusion community to date (it is not easy, conceptually or technically), yet it represents a major theme for design integration. This section outlines some of the ideas, raises some questions and perhaps triggers new thinking and approaches. Handling uncertainty in design is not new of course—organisations such as NASA include it in approaches to design [84–86] even if the nature of the uncertainties is different. Managing uncertainties quantitatively is a major task—the integration tools of section 8 can be a key part of the strategy.

7.2. Types of uncertainty, their implications and management

In the uncertainty community there are two main classes of uncertainty: aleatory and epistemic. The definitions vary, but can be approximately written:

- ‘Aleatory’—non-deterministic phenomena, statistical or irreducible uncertainty (such as measurements bounded by photon statistics)
- ‘Epistemic’—gaps in understanding (data, models, environment)—this includes the ‘unknown unknowns’, and potentially unforeseeable behaviours. A first principles model may exist, but if the time and/or computational resources have not been available to use it, an uncertainty is left.

The situation is made more intricate due to the deeply coupled systems in fusion which can also couple the different types of uncertainties. It is often not possible, or useful, to allocate an uncertainty to one category or another uniquely, as many uncertainties have a mix of both, and when estimating the overall performance uncertainty both types appear: e.g. there may be an epistemic uncertainty in how data with aleatory uncertainty is propagated. It is, however, still important to recognise the two general categories as they profoundly affect the approaches to handling uncertainty. If an uncertainty can be defined as aleatory, or statistical, that allows powerful statistical methods to be employed to combine and propagate uncertainties through a system. So, even if an uncertainty may not be irreducible (e.g. it can be reduced with further work), and the allocation of a particular statistical distribution (e.g. Gaussian or Poisson) may not be justified, it allows progress to be made, as long as the assumptions are explicit and can be tracked and changed later if needed.

One of the big questions is knowing how to handle the epistemic uncertainties, in particular how to proceed when the models are incomplete (or non-existent for some phenomena). This is a rather different situation from using data to select models. For example, in an existing tokamak one has data (with uncertainties) on profiles and often turbulence which can be used as input to transport models. These can in turn be used to predict the global energy confinement time which can be measured largely independently. Thus one can, perhaps by using Bayesian techniques, choose which transport model(s) best represent the full data set. The situation for a power plant design is rather different: there is often very little input data

(besides atomic and nuclear data to feed plasma and materials models) and of course no data on the overall performance. While the models are incomplete (which will ‘always’ be the case to some extent, and *a fortiori* when early decisions are needed, as for most of the DEMO-like programmes), how does one choose which of these incomplete models to trust for predictions and uncertainty propagation? This is at the heart of the problem.

There are at least three other sources of variation and uncertainty important in assessing the suitability of the design:

- Response to disturbances—unplanned behaviour of control actuators (especially heating and fuelling systems), confinement transitions, impurity fragments etc
- Impact of uncertainty in initial or boundary conditions
- Uncertainty in the stakeholder requirements over time.

The first two should be tractable via a suitable set of models, while noting that incomplete models may not give trustworthy results. Tools developed for ‘aleatory’ uncertainty propagation can probably be used, and such frameworks also allow different models or sub-models to be compared (again they do not necessarily determine which models are trustworthy).

The final one is not addressed explicitly here, rather it is assumed to be handled by margins and agile design, and the margins and agility need to bear both internal and external stakeholders in mind (the position of external stakeholders and investors is likely to be influenced by the understanding and confidence in the internal community). We now look at some examples and possible approaches.

72.1. Main device parameters. One approach to handle uncertainties is to use scaling relations (empirical or theoretical) and then vary the input parameters, such as plasma density, hypothetical confinement multipliers (e.g. H-factor), elongation etc and propagate through a systems code optimiser to see the impact on the design [87, 88]. However, as mentioned earlier, caution is needed on the use of large multi-source databases for choosing an optimum design of a particular device, and of course use of multipliers on scalings beyond the statistical spread has weak rigour. The use of scalings to propagate uncertainties (e.g. within systems codes) also needs to be examined for rigour and consistency, as it could take the scaling beyond its range of validity, and use of 0D scalings may unwittingly propagate assumptions about plasma profiles which are not supported. For example, pedestal density, which affects the total plasma energy content and thus the predicted fusion power, may result from neutral penetration which is minimal in a DEMO-class device. An alternative is to explore the impact of a range of assumptions in the models (e.g. different pedestals for the same core transport).

72.2. Response to disturbances and initial/boundary conditions. If/when detailed integrated first principles models are available, then sophisticated uncertainty quantification (UQ) is possible—but with a very different character,

since the model is in principle accurate¹². The uncertainty now arises from changes in initial or boundary conditions: at the design stage (impact of different location of pellets, gas valves, heating systems, wall shape etc), once the design has been finalised (variations from the concept design which may affect the achievable performance range) and in operation (response to disturbances). All of these should start to be assessed from early in the design, even with uncertain models. The tool of preference would be a UQ-enabled *flight simulator* with interfaces to the technology, engineering and systems (see below)—sometimes UQ-enabled simulators are termed ‘emulators’ in the Uncertainty community.

72.3. Technology and design implications. From these two areas are likely to include:

- Range of operating points/zones to meet the devices goals, during the design phase at least
- Changes in the fusion and exhaust power due to disturbances or revised physics
- Necessary control actuator reliability, redundancy and speed of switch-in of reserves
- In-vessel coils for divertor sweeping in case of reattachment [52]
- Diagnostic/observer precision required (see [78] for discussion of diagnostics for European DEMO).

72.4. Exploiting uncertainties. The UQ tools for aleatory uncertainties can also be a powerful way to exploit uncertainties—outlying points can be explored for enhanced performance.

72.5. Practicalities of UQ. The computational challenge is substantial—large numbers of runs of high fidelity codes (probably to generate reduced fidelity-conserving models or data fits) are needed and the resulting big datasets must be analysed effectively. How many runs is a subject of discussion and can vary a lot between model types and between the techniques to create usable emulators from the experimental or modelling data. There is a rule of thumb, that for the right type of problem, about ten times the number of runs as input parameters are required to build a Gaussian Process Emulator.

72.6. Epistemic uncertainties—examples. The nature and path to an ELM-free pedestal is not yet known. How high is it? Does it only exist in certain plasma shapes, or in certain parameter ranges? Managing the uncertainty might need a larger device (e.g. if the core transport must exploit saturated ITG turbulence and be ‘stiff’ for the fusion power to be controllable), or improved core plasma confinement (with its own uncertainties), if attainable. If a larger device is chosen and the no-ELM pedestal turns out to be of a similar height to

¹² The ideal model may never be reached, but as the model becomes more first-principles, then uncertainties resulting from fits to uncertain experimental data (e.g. via fitting parameters with statistical variations) start to disappear.

the standard model (close to the peeling-ballooning limit) and one needs stiff core transport for control and reproducibility reasons, the total fusion power may be higher, affecting the rest of the plant.

7.2.7. Quantifying epistemic uncertainties. These are intrinsically difficult to quantify—by definition we do not have reliable models (or, as above, might not have been able to run them even if they exist). One approach (which has been proposed for weather systems [89]) is to use ‘storylines’, or ‘what-if’ scenarios. To quote [89]: ‘Moreover, the recognition that epistemic uncertainties are deterministic removes the impulse to provide probabilities, which can give the illusion of objectivity and thereby reduce transparency. Instead, epistemic uncertainty can be represented through a discrete set of (multiple) storylines—physically self-consistent, plausible pathways, with no probability attached [...]. Rather than asking what will happen (as in the traditional, scenario-driven approach), which we may not be able to answer with any confidence, storylines allow us to ask what would be the effect of particular interventions—e.g. different climate forcing scenarios, or different adaptation measures—across a range of plausible futures.’ i.e. it is about making assumptions and exploring the consequences—in this case the integrated models are more ‘assumption integrators’ than true predictors. A next step would be to aim for a design that is optimised for ‘insensitivity’ to the epistemic uncertainties, or ‘insensitivity to the story’. Whether or not this optimisation is adopted, the approach allows one to see the consequences of certain assumptions either to the design or to the final device if the situation materialises and the machine was not explicitly designed to accommodate it. This emphasises the central importance of adaptable integrated models and flight simulators, while being very explicit about what they are and are not.

To close we give some examples of epistemic uncertainties in an integrated design: some are gaps in theory, data and fundamental understanding (and thus credible models), others may be gaps in the application of the knowledge due to algorithm or CPU limitations.

- (a) The pedestal uncertainty outlined above, and the non-linear coupling/interplay with the core plasma
- (b) Core transport magnitude and behaviour in the new regimes (affects P_{fus} and controllability)—likely to have modest uncertainty for the current European DEMO scenarios and e.g. ARC, but could have large uncertainties for more ambitious plasma concepts
- (c) Materials lifetimes and property changes under irradiation which would affect the allowed neutron fluence between component replacement. This may also be affected by the waste management strategy and regulatory regime.

7.3. Integrating in the presence of uncertainty

Possible approaches could be:

- Develop several engineering-consistent concepts for different assumed plasma scenarios and physics and decide which

to construct when the uncertainties of each are reduced (e.g. more data, better models)

- A single design that can accommodate more than one plasma scenario (this is the flexi-DEMO approach [81])—of course one needs to know with sufficient accuracy what device parameters each scenario needs
- Pick a concept based on best judgement of researchers or stakeholders (even if not supported by strong evidence) and make a full design based on a set of assumptions and then rely on operational optimisation, such as has been demonstrated on several devices, e.g. JET’s performance recovery after installation of the ITER-like metal wall, or modifications *in-situ* via an upgradable design (e.g. highly modular) and very flexible remote maintenance and replacement systems.
- Choose concepts that eliminate some plasma and technology elements which have high uncertainty or particularly large challenges, e.g. by different architectures, even if in exchange for others which are innovative and untested. For example ARC [27] aims to use plasma parameters in the existing experiments regime, and transfer some of the challenge to the attainment of high field magnets. STs [90–92] innovate to address remote maintenance challenges, some stellarator options focus on less challenging superconducting magnet technology at expense of lower field [93, 94].
- Estimating and assessing the integrated uncertainty: as indicated above this is not mechanistic due to the combination of the different types of uncertainty (especially knowledge gaps), as well as the technical challenges of estimating. One approach is to create ‘storylines’ based on collections of assumptions (which would reduce over time) and then calculate the integrated plant performance bands for each. Determining the acceptability may not be amenable to quantitative criteria.

The choices will depend on the community and stakeholders.

8. Integration tools

We have discussed many examples and aspects of integration above, showing the importance of organised tools to handle the complexity of the interfaces and optimisation, while at the same time indicating the danger of trusting algorithms when the models are incomplete and there are knowledge gaps (epistemic uncertainties).

There are presently a few classes of computational tools integrating the physics, technology and engineering: systems codes, a new class of engineering scoping tools and early-stage flight simulators. These are supplemented by other more localised tools, for example integrated plasma models such as JIN-TRAC (used in [54]), integrated turbulence models [95], parametric engineering models for optimising breeding blankets including neutronics [96] and materials models that translate the theory of radiation effects into their engineering impact, e.g. [97]. When a plant is in operation, the use of digital twins of various kinds could be powerful—for example it might

be possible to assess, given the detailed history of loading and thermal fatigue, the state of the surface of the PFCs with respect to cracking and the potential to release fragments of tungsten of various sizes (one of the drivers and challenges for the heating capability for the European DEMO is the ability to compensate the cooling from such fragments [80]).

Traditionally systems codes [28, 29, 98, 99] have been used for multi-variable optimisation using a set of simplified physics (often 0D) and engineering rules as the first step in a design process. More recently the realisation of the importance of higher order and more detailed effects (such as ELMs, some limitations in 0D scalings and the cross-cutting integration challenges of the European DEMO) has triggered a combination of greater sophistication in the systems codes, increased emphasis on more advanced and complete models, more questioning of (a) the outputs of systems codes including the value of their optimisations, and (b) the general approach of starting with low fidelity and moving to high fidelity (as well as missing constraints and problems it can also miss opportunities and new regimes). However systems codes provide a powerful way to integrate the whole plant quickly—including the power conversion and cost models—to show the interfaces and tensions explicitly.

The optimisation of the TF and PF coil system explored in the alternative divertor configuration above, where unconventional shapes of TF coils are considered shows the benefit of fast magnet optimisation codes that can be linked to first stage finite element mechanical engineering analysis and optimisation. The BLUEPRINT code originated from this need and is being expanded to a wider capability (it can generate a simplified 3D CAD model in a few seconds) [36, 37], and such models could have more technology added [100, 101]. Caution is still needed given the internal complexity of a TF coil for example.

Finally, and most importantly here, is the so-called flight simulator. In simple terms the aim is to have a time-dependent model that simulates the whole plasma pulse including the control systems, diagnostics, and external constraints such as power load to the PFCs. As well as exhibiting the general behaviour, perhaps the most interesting aspects are control, and showing how the impact of various uncertainties can be estimated and managed (e.g. disturbances from partial system failures, the effect of different transport models and assumptions)—recall the point above about the importance of ‘UQ-enabled’ simulators (or emulators). It is expected that flight simulators will be vital to qualifying designs, optimising and minimising plasma commissioning (see above), as well as developing the concept in the first place. Basic versions already exist for the core plasma in ASDEX Upgrade and the European DEMO [102–104], while a full version is some way away and will need novel computational techniques to allow reduced high fidelity simulations (e.g. non-linear turbulence) that run fast enough.

Together these have the potential to allow rapid substantiated design. In particular they may allow substantial changes late in the sequence—one of the ways of handling epistemic and other uncertainties and preventing too tight a design integration too early (restricting innovation and even solutions).

However the manufacturing supply chain may not be able to react quickly to significant changes; integration with the manufacturers is important.

These tools can serve another important function: knowledge integration, preservation and transfer. In principle they should use state of the art knowledge and their documentation encapsulate the reasons behind the choices of models and theory used. This can help avoid ‘unknown knowns’—knowledge of which newcomers are not aware.

9. Summary and implications

In this brief tour of a wide topic, we have described why it is important to take early consideration of integration across the whole plant (physics of plasma and materials, engineering and multiple space and time scales) and its lifecycle (over a single pulse and over the plant lifetime including commissioning). The journey has taken us into some quite deep themes about fusion plant design, has reinforced the importance of designing with the end point in mind, to construct R&D based on looking at the end state in detail, and to complement the usual discovery-focused nature of research programmes where the research direction tends to change with the discoveries. We have listed many integration aspects and have delved more deeply into three examples (exhaust, tritium management, plasma scenario) where such integration illuminates the situation, revealing new solutions as well as new challenges. The implications of integration can be decisive in modifying and eliminating existing options, i.e. there are far-reaching implications on and of the plasma scenario.

Integration factors such as recirculating power, divertor configuration, PFCs, ELMs etc are well known. However, there are many others, some perhaps initially surprising, for example:

Controllability and transport: it may be more important to consider controllability than a particular fusion power and aim for transport characteristics (e.g. profile clamping) that suit control rather than aim for maximum confinement time.

Plasma Commissioning: the final plasma will have dominant alpha-particle heating and be to some extent self-organising, the heating related to the transport. Consideration is needed about whether and how it can be commissioned in a non-active phase (e.g. in protium).

Power range and blanket design: if the natural (e.g. controllable) plasma has higher fusion power than planned, then the blanket coolant and heat transfer systems, and the balance of plant (heat exchangers etc) will need margin.

Energy stores: energy stores may be needed for control excursions and for plasma ramp-up and ramp down (when there is little or no fusion power) as well as mitigating fatiguing of the balance of plant—filling these adds to the recirculating power, which always has to be minimised.

Tritium buffer and fuelling options: the necessary tritium buffer depends on lag times in the fast internal fuel loop as the D-T supply to the plasma is ramped up, and on the retention of tritium in materials as they suffer radiation damage. Further it

may be impractical to vary the D:T fuel ratio on a short timescale for power control.

Power density and remote maintenance: higher fusion power density (smaller devices) has implications on materials lifetime and the speed of robotic maintenance. These all feed back into the plasma optimisation and concept selection.

An emerging element is integration in the presence of uncertainty in the plasma and technology performance. Uncertainty can provide space and focus for innovation and improvement, but tends to make specific engineering design more difficult. Developing power plants at a feasible cost on a short timescale requires large extrapolations with limited experimental data—very different from many other major technology and engineering endeavours where progress is generally made by incremental experiments, devices and prototypes usually with relatively modest advances each time (in the grand scheme). These large extrapolations have associated uncertainties, especially knowledge gaps (termed epistemic uncertainties) in some important areas—how many and how significant depends on the global and local concept chosen. There are other uncertainties such as the impact of uncertainties in plasma boundary conditions or constraints; under what conditions, after how long and in what way components may fail; the impact of failures and performance variations in the control systems. All the various uncertainties must be managed in the presence of time pressure from stakeholders (e.g. driven by the urgency of tackling climate change) whose requirements may potentially evolve over time, affecting the design. The issue of uncertainties leads us to new conceptual, intellectual and practical challenges: there is a range of ideas and methodologies for handling uncertainty, but most have to be developed and tested for their feasibility and timeliness. Handling uncertainty will probably have to become an intrinsic element of plasma and engineering optimisation and integration tools: estimating, propagating, and accommodating the consequences via margin and innovation. It will require ingenuity as well as large scale computation.

There are two main ways to address the uncertainties and fortunately these align with the tools to develop integrated designs in the first place. One way is to develop and use reduced (i.e. fast) high fidelity plasma (and component) models, informed by experimental data from relevant facilities (noting that the test regimes may not be representative of the final environment). These allow the impact of uncertainties to be estimated and the design optimised, and tools to identify modified or alternative plasma scenarios. Another way is to develop concepts with engineering margin and/or development strategies that can accommodate significant performance uncertainties and alternative scenarios. Each of these should handle higher or lower risk approaches depending on stakeholder appetite, e.g. accepting less mature but potentially higher performance technologies or plasma ingredients/concepts.

A major element of the first path will be so-called plasma flight simulators that combines physics models across the whole plasma with technology constraints/behaviours, especially control observers and actuators. These simulators only exist in simplified forms today and their development is likely

to be a major theme in research programmes in future—indeed experiments could be designed in terms of their contributions to such simulators. They will need to exploit the transformational computational developments in algorithms and hardware (e.g. the exascale). The simulators can also help with another integration challenge—knowledge repositories and transfer across generations of researchers, e.g. via rich annotation and literature references.

For the second path, two approaches offer promise: device concepts that can accommodate more than one plasma scenario (e.g. pulsed and steady state, or lower and higher fusion power) and advanced digital design tools that allow fast high fidelity engineering models (e.g. rapid generation of high resolution models and fast yet detailed mechanical, thermal and thermohydraulic analysis) which can allow late yet rapid and comprehensive design changes. Some of these are the engineering and technology analogues of the plasma flight simulator. Again these can be enabled in part by transformative advances in computation.

Integration of the plasma with the rest of the plant reveals many aspects and challenges often not considered within the specialised research programmes that have underpinned the development of fusion so far. ITER has brought many to light, and programmes looking further ahead, notably EUROfusion's DEMO, are revealing a new tranche. We have picked a few examples in this paper to try to show how the constraints and challenges from integration are not insurmountable, and that consideration of integration early on can uncover some interesting and attractive new approaches.

Data availability statement

No new data were created or analysed in this study.

Acknowledgments

The authors would like to acknowledge important discussion with Mohamad Abdallah, Mike Gorley, Jonathan Graves, James Harrison, Andrew Kirk and Mattia Siccino, amongst many others over recent years. This work has been funded by the EPSRC Energy Programme [Grant Number EP/T012250/1]. This work has been carried out within the framework of the EUROfusion Consortium and has received funding from the Euratom research and training programme 2014–2018 and 2019–2020 under Grant Agreement No. 633053. The views and opinions expressed herein do not necessarily reflect those of the European Commission.

ORCID iDs

A W Morris  <https://orcid.org/0000-0001-7281-5717>
F Militello  <https://orcid.org/0000-0002-8034-4756>
E Surrey  <https://orcid.org/0000-0003-3093-9556>
H R Wilson  <https://orcid.org/0000-0003-3333-7470>

References

- [1] Federici G *et al* 2014 Overview of EU DEMO design and R&D activities *Fusion Eng. Des.* **89** 882
- [2] Federici G *et al* 2019 Overview of the DEMO staged design approach in Europe *Nucl. Fusion* **59** 066013
- [3] Various Authors 2022 Special Issue on European Programme towards DEMO: Outcome of the Pre-Conceptual Design Phase *Fusion Eng. Des.* (available at: www.sciencedirect.com/journal/fusion-engineering-and-design/special-issue/10RRZQ6LW4H)
- [4] ITER (available at: www.iter.org/)
- [5] Bigot B *et al* 2022 Preparation for assembly and commissioning of ITER *Nucl. Fusion* **62** 042001
- [6] EUROfusion 2018 European research roadmap to the realisation of fusion energy (available at: www.euro-fusion.org/eurofusion/roadmap/)
- [7] JT-60SA Advanced Superconducting Tokamak (available at: www.jt60sa.org/)
- [8] Dinklage A *et al* 2018 Magnetic configuration effects on the Wendelstein 7-X stellarator *Nat. Phys.* **14** 855
- [9] Klinger T *et al* 2019 Overview of first Wendelstein 7-X high-performance operation *Nucl. Fusion* **59** 112004
- [10] Milnes J, Ayed N B, Dhalla F, Fishpool G, Hill J, Katramados I, Martin R, Naylor G, O’Gorman T and Scannell R 2015 MAST upgrade—construction status *Fusion Eng. Des.* **96–97** 42
- [11] Morris A W 2018 MAST upgrade divertor facility: a test bed for novel divertor solutions *IEEE Trans. Plasma Sci.* **46** 1217
- [12] Harrison J R *et al* 2022 Initial demonstration of enhanced divertor heat flux mitigation and detachment access in the MAST Upgrade Super-X divertor configuration *Submitted to Phys. Rev. Lett.*
- [13] Menard J E *et al* 2012 Overview of the physics and engineering design of NSTX Upgrade *Nucl. Fusion* **52** 083015
- [14] Lee G S *et al* 2000 The KSTAR project: an advanced steady state superconducting tokamak experiment *Nucl. Fusion* **40** 575
- [15] Wu S the EAST Team 2007 An overview of the EAST project *Fusion Eng. Des.* **82** 463
- [16] Albanese R *et al* 2017 The DTT proposal. A tokamak facility to address exhaust challenges for DEMO: introduction and executive summary *Fusion Eng. Des.* **122** 274
- [17] Creely A J *et al* 2020 Overview of the SPARC tokamak *J. Plasma Phys.* **86** 865860502
- [18] Najmabadi F, Conn R and the ARIES team 1991 The ARIES-I tokamak reactor study *Fusion Technol.* **19** 783
- [19] Kessel C E *et al* 2015 The ARIES advanced and conservative tokamak power plant study *Fusion Sci. Technol.* **67** 1
- [20] Tobita K *et al* 2019 Japan’s efforts to develop the concept of JA DEMO during the past decade *Fusion Sci. Technol.* **75** 372
- [21] Ishii I *et al* 2021 R&D activities for fusion DEMO in the QST Rokkasho Fusion Institute *Fusion Sci. Technol.* **77** 532–48
- [22] Cho A, Kwon J-M, Chung H-K, Kim J, Kang J, Choi W-J and Lee E S 2022 A planning study for virtual DEMO development in Korea *Fusion Eng. Des.* **176** 113026
- [23] Zhuang G *et al* 2019 Progress of the CFETR design *Nucl. Fusion* **59** 112010
- [24] Gliss C, Ciattaglia S, Korn W and Moscato I 2018 Initial layout of DEMO buildings and configuration of the main plant systems *Fusion Eng. Design* **136** 534
- [25] Bachmann C *et al* 2021 Containment structures and port configurations *Fusion Eng. Des.* **174** 112966
- [26] Bachmann C *et al* 2020 Key design integration issues addressed in the EU DEMO pre-concept design phase *Fusion Eng. Design* **156** 111595
- [27] 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
- [28] 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
- [29] 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
- [30] ITER Physics Basis Editors *et al* 1999 ITER physics basis *Nucl. Fusion* **39** 2137
- [31] Shimada M *et al* 2007 Progress in the ITER physics basis *Nucl. Fusion* **47** S1
- [32] Siccino M, Fable E, Angioni C, Saarelma S, Scarabosio A and Zohm H 2018 Impact of an integrated core/SOL description on the R and B_T optimization of tokamak fusion reactors *Nucl. Fusion* **58** 016032
- [33] Sips A A C *et al* 2007 The performance of improved H-modes at ASDEX upgrade and projection to ITER *Nucl. Fusion* **47** 1485
- [34] Schweinzer J *et al* 2011 Confinement of ‘improved H-modes’ in the all-tungsten ASDEX Upgrade with nitrogen seeding *Nucl. Fusion* **51** 113003
- [35] Mailloux J *et al* 2022 Overview of JET results for optimising ITER operation *Nucl. Fusion* (<https://doi.org/10.1088/1741-4326/ac47b4>)
- [36] Coleman M and McIntosh S 2019 BLUEPRINT: a novel approach to fusion reactor design *Fusion Eng. Des.* **139** 26
- [37] Coleman M and McIntosh S 2020 The design and optimisation of tokamak poloidal field systems in the BLUEPRINT framework *Fusion Eng. Des.* **154** 111544
- [38] Albanese R, Ambrosino R, Castaldo A and Loschiavo V P 2018 Optimization of the PF coil system in axisymmetric fusion devices *Fusion Eng. Des.* **133** 163
- [39] Maviglia F, Albanese R, Ambrosino R, Bachmann C, Federici G and Villone F 2019 Optimization of DEMO geometry and disruption location prediction *Fusion Eng. Des.* **146** 967
- [40] Wenninger R, Kemp R, Maviglia F and Zohm H 2015 DEMO exhaust challenges beyond ITER *42nd EPS Conference on Plasma Physics Lisbon* p P4.110 (available at: <http://ocs.ciemat.es/EPS2015PAP/pdf/P4.110.pdf>)
- [41] Wenninger R *et al* 2017 The DEMO wall load challenge *Nucl. Fusion* **57** 046002
- [42] Zohm H, Militello F, Morgan T W, Morris W, Reimerdes H and Siccino M 2021 The EU strategy for solving the DEMO exhaust problem *Fusion Eng. Des.* **166** 112307
- [43] Eich T *et al* 2011 Inter-ELM power decay length for JET and ASDEX Upgrade: measurement and comparison with heuristic drift-based model *Phys. Rev. Lett.* **107** 215001
- [44] Eich T *et al* 2013 Scaling of the tokamak near the scrape-off layer H-mode power width and implications for ITER *Nucl. Fusion* **53** 093031
- [45] Hasegawa A, Tanno T, Nogami S and Satou M 2011 Property change mechanism in tungsten under neutron irradiation in various reactors *J. Nucl. Mater.* **417** 491
- [46] Mason D R, Reza A, Granberg F and Hofmann F 2021 Estimate for thermal diffusivity in highly irradiated tungsten using molecular dynamics simulation *Phys. Rev. Mater.* **5** 125407
- [47] Ariola M, Pironti A, Ambrosino R, Mattei M, Biel W and Franke T 2019 Simulation of magnetic control of the plasma shape on the DEMO tokamak *Fusion Eng. Des.* **146** 728
- [48] Loarte A, Koechl F, Leyland M J, Polevoi A, Beurskens M, Parail V, Nunes I, Saibene G R and Sartori R I A 2014 Evolution of plasma parameters in the termination phase

- of high confinement H-modes at JET and implications for ITER *Nucl. Fusion* **54** 123014
- [49] Fuchert G *et al* 2020 Increasing the density in Wendelstein 7-X: benefits and limitations *Nucl. Fusion* **60** 036020
- [50] Lehnen M *et al* 2015 Disruptions in ITER and strategies for their control and mitigation *J. Nucl. Mater.* **463** 39
- [51] de Vries P C, Johnson M F, Alper B, Buratti P, Hender T C, Koslowski H R and Riccardo V 2011 Survey of disruption causes at JET *Nucl. Fusion* **51** 053018
- [52] Maviglia F *et al* 2022 Integrated design strategy for EU-DEMO first wall protection from plasma transients *Fusion Eng. Des.* **177** 113067
- [53] Bernert M *et al* 2021 X-point radiation, its control and an ELM suppressed radiating regime at the ASDEX Upgrade tokamak *Nucl. Fusion* **61** 024001
- [54] Koechl F *et al* 2020 Evaluation of fuelling requirements for core density and divertor heat load control in non-stationary phases of the ITER DT 15 MA baseline scenario *Nucl. Fusion* **60** 066015
- [55] Remeirdes H *et al* 2020 Assessment of alternative divertor configurations as an exhaust solution for DEMO *Nucl. Fusion* **60** 066030
- [56] Militello F *et al* 2021 Preliminary analysis of alternative divertors for DEMO *Nucl. Mater. Energy* **26** 100908
- [57] Militello F *et al* 2021 An assessment of alternative divertors for the European DEMO 28th IAEA Fusion Energy Conference Virtual 10-15 May 2021 TH/P4-9 (available at: <https://nucleus.iaea.org/sites/fusionportal/Shared%20Documents/FEC%202020/fec2020-preprints/preprint1083.pdf>)
- [58] Kembleton R, Siccino M, Maviglia F and Militello F 2022 Benefits and challenges of advanced divertor configurations in DEMO *Fusion Eng. Des.* **179** 113120
- [59] Barabaschi P Private communication 2019 Uniformity of thick forgings
- [60] Valanju P M, Kotschenreuther M, Mahajan S M and Canik J 2009 Super-X divertors and high power density fusion devices *Phys. Plasmas* **16** 056110
- [61] Xiang L, Militello F, Mouylton D, Subba F, Aho-Mantila L, Coster D, Wensing M, Lunt T, Wischmeier M and Reimerdes H 2021 The operational space for divertor power exhaust in DEMO with a super-X divertor *Nucl. Fusion* **61** 076007
- [62] Lipschultz B, Parra F I and Hutchinson I H 2016 Sensitivity of detachment extent to magnetic configuration and external parameters *Nucl. Fusion* **56** 056007
- [63] Umansky M V, LaBombard B, Brunner D, Golfinopoulos T, Kuang A Q, Rensink M E, Terry J L, Wigram M and Whyte D G 2020 Study of passively stable, fully detached divertor plasma regimes attained in innovative long-legged divertor configurations *Nucl. Fusion* **60** 016004
- [64] Fasoli A *et al* 2020 TCX heating and divertor upgrades *Nucl. Fusion* **60** 016019
- [65] Herrmann A *et al* 2017 An optimized upper divertor with divertor-coils to study enhanced divertor configurations in ASDEX Upgrade *Fusion Eng. Des.* **123** 508
- [66] Guo H Y, Sang C F, Stangeby P C, Lao L L, Taylor T S and Thomas D M 2017 Small angle slot divertor concept for long pulse advanced tokamaks *Nucl. Fusion* **57** 044001
- [67] Lang P T, Day C, Fable E, Igitkhanov Y, Köchl F, Mooney R, Pegourie B, Ploeckl B, Wenninger R and Zohm H 2015 Considerations on the DEMO pellet fuelling system *Fusion Eng. Des.* **96–97** 123
- [68] Lang P T, Cismondi F, Day C, Fable E, Frattolillo A, Gliss C, Janky F, Pégourié B and Ploeckl B 2020 Optimizing the EU-DEMO pellet fuelling scheme *Fusion Eng. Des.* **156** 111591
- [69] Abdou M, Riva M, Ying A, Day C, Loarte A, Baylor L, Humrickhouse P, Fuerst T F and Cho S 2021 Physics and technology considerations for the deuterium–tritium fuel cycle and conditions for tritium fuel self sufficiency *Nucl. Fusion* **61** 013001
- [70] Derlet P M and Dudarev S L 2020 Microscopic structure of a heavily irradiated material *Phys. Rev. Mater.* **4** 023605
- [71] Mason D R, Granberg F, Boleining M, Schwarz-Selinger T, Nordlund K and Dudarev S L 2021 Parameter-free quantitative simulation of high-dose microstructure and hydrogen retention in ion-irradiated tungsten *Phys. Rev. Mater.* **5** 095403
- [72] Hollingsworth A *et al* 2020 Comparative study of deuterium retention in irradiated Eurofer and Fe–Cr from a new ion implantation materials facility *Nucl. Fusion* **60** 016024
- [73] Hollingsworth A *et al* 2022 Comparative study of deuterium retention and vacancy content of self-ion irradiated tungsten *J. Nucl. Mater.* **558** 133373
- [74] Wang J *et al* 2021 Deuterium retention in W and binary W alloys irradiated with high energy Fe ions *J. Nucl. Mater.* **545** 152749
- [75] Day C, Butler B, Giegerich T, Ploeckl B and Varoutis S 2019 A smart three-loop fuel cycle architecture for DEMO *Fusion Eng. Des.* **146** 2462
- [76] Giegerich T and Day C 2014 The KALPUREX-process—a new vacuum pumping process for exhaust gases in fusion power plants *Fusion Eng. Des.* **89** 1476
- [77] Hartl T *et al* 2022 Design and feasibility of a pumping concept based on tritium direct recycling *Fusion Eng. Des.* **174** 112969
- [78] Biel W *et al* 2019 Diagnostics for plasma control—from ITER to DEMO *Fusion Eng. Des.* **146** 465
- [79] Saarelma S, Challis C D, Garzotti L L, Frassinetti L, Maggi C F, Romanelli M, Stokes C *et al* 2018 Integrated modelling of H-mode pedestal and confinement in JET-ILW *Plasma Phys. Control. Fusion* **60** 014042
- [80] Siccino M *et al* 2022 Impact of the plasma operation on the technical requirements in EU-DEMO *Fusion Eng. Des.* **179** 113123
- [81] Zohm H, Träuble F, Biel W, Fable E, Kemp R, Lux H, Siccino M and Wenninger R 2017 A stepladder approach to a tokamak fusion power plant *Nucl. Fusion* **57** 086002
- [82] Siccino M, Graves J P, Kembleton R, Lux H, Maviglia F, Morris A W, Morris J and Zohm H 2022 Development of the plasma scenario for EU-DEMO: status and plans *Fusion Eng. Des.* **176** 113047
- [83] Beidler C D *et al* 2021 Demonstration of reduced neoclassical energy transport in Wendelstein 7-X *Nature* **596** 221
- [84] Mauery T *et al* 2021 A guide for aircraft certification by analysis NASA, NASA/CR-20210015404
- [85] Liu X, Furrer D, Kusters J and Holmes J 2018 Vision 2040: a roadmap for integrated, multiscale modeling and simulation of materials and systems NASA/CR—2018-219771
- [86] Novack S D, Rogers J, Al Hassan M and Hark F 2016 Characterizing epistemic uncertainty for launch vehicle designs 8th IAASS Conference “Safety First, Safety for All” Melbourne, FL 2016 NASA Document 20160006978 (available at: <https://ntrs.nasa.gov/api/citations/20160006978/downloads/20160006978.pdf>)
- [87] Lux H, Kemp R, Wenninger R, Biel W, Federici G, Morris W and Zohm H 2017 Uncertainties in power plant design point evaluations *Fusion Eng. Des.* **123** 63
- [88] Lux H, Siccino M, Biel W, Federici G, Kembleton R, Morris A W, Patelli E and Zohm H 2019 Implications of uncertainties on European DEMO design *Nucl. Fusion* **59** 066012
- [89] Shepherd T G 2019 Storyline approach to the construction of regional climate change information *Proc. R. Soc. A* **475** 20190013

- [90] Menard J E *et al* 2016 Fusion nuclear science facilities and pilot plants based on the spherical tokamak *Nucl. Fusion* **56** 106023
- [91] Menard J E 2019 Compact steady-state tokamak performance dependence on magnet and core physics limits *Phil. Trans. R. Soc. A* **377** 21070440
- [92] UKAEA STEP: Spherical Tokamak for Energy Production (available at: <https://step.ukaea.uk/>)
- [93] Helander P, Drevlak M, Zarnstorff M and Cowley S C 2020 Stellarators with permanent magnets *Phys. Rev Lett.* **124** 095001
- [94] Zarnstorff M *et al* 2021 Simpler optimized stellarators using permanent magnets *47th EPS Conf. on Plasma Physics* p P4.1055 (available at: <http://ocs.ciemat.es/EPS2021PAP/pdf/P4.1055.pdf>)
- [95] Bhattacharjee A *et al* 2021 Accelerating magnetically confined fusion through advancements in edge turbulence modeling and its integration in a whole device model *28th IAEA Fusion Energy Conference Virtual 10-15 May 2021 OV/4-1* (available at: <https://nucleus.iaea.org/sites/fusionportal/Shared%20Documents/FEC%202020/fec2020-preprints/preprint1357.pdf>)
- [96] Shimwell J *et al* 2019 Multiphysics analysis with CAD-based parametric breeding blanket creation for rapid design iteration *Nucl. Fusion* **59** 046019
- [97] Dudarev S L, Mason D R, Tarleton E, Ma P-W and Sand A E 2018 A multi-scale model for stresses, strains and swelling of reactor components under irradiation *Nucl. Fusion* **58** 126002
- [98] Wenninger R *et al* 2017 The physics and technology basis entering European system code studies for DEMO *Nucl. Fusion* **57** 016011
- [99] Reux C *et al* 2018 DEMO design using the SYCOMORE system code: influence of technological constraints on the reactor performances *Fusion Eng. Des.* **136** 1572
- [100] Franza F, Boccaccini L V, Fischer U, Gade P V and Heller R 2015 On the implementation of new technology modules for fusion reactor systems codes *Fusion Eng. Des.* **98–99** 1767
- [101] Franza F 2019 Development and validation of a computational tool for fusion reactors' system analysis *PhD Thesis* Karlsruhe Institute of Technology
- [102] Fable E, Janky F, Treutterer W, Englberger M, Schramm R, Muraca M *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
- [103] Janky F, Fable E, Treutterer W and Zohm H the EUROfusion-IM Team 2017 Simulation of burn control for DEMO using ASTRA coupled with simulink *Fusion Eng. Des.* **123** 555
- [104] Janky F, Fable E, Treutterer W, Gomez Ortiz I and Kudlacek O the ASDEX Upgrade Team 2019 ASDEX upgrade flight simulator development *Fusion Eng. Des.* **146** 1926