PAPER • OPEN ACCESS

Plasma control for the step prototype power plant

To cite this article: M. Lennholm *et al* 2024 *Nucl. Fusion* **64** 096036

OPEN ACCESS
IOP Publishing | International Atomic Energy Agency
Nuclear Fusion
Nucl. Fusion **64** (2024) 096036 (20pp)
https://doi.org/10.1088/1741-4326/ad6012

Plasma control for the step prototype power plant

M. Lennholm[1,*], S. Aleiferis[1], S. Bakes[1], O.P. Bardsley[1], M. van Berkel[2], F.J. Casson[1], F. Chaudry[1], N.J. Conway[1], T.C. Hender[1], S.S. Henderson[1], A. Hudoba[1], B. Kool[2,3], M. Lafferty[1], H. Meyer[1], J. Mitchell[1], A. Mitra[1], R. Osawa[1], R. Otin[1], A. Parrott[1], T. Thompson[1], G. Xia[1] and the STEP Team[1]

[1] United Kingdom Atomic Energy Authority, Culham Campus, Abingdon, Oxon, OX14 3DB, United Kingdom of Great Britain and Northern Ireland
[2] DIFFER—Dutch Institute for Fundamental Energy Research, De Zaale 20, 5612 AJ Eindhoven, Netherlands
[3] Department of Mechanical Engineering, Control Systems Technology Group, Eindhoven University of Technology, Eindhoven, Netherlands

E-mail: morten.lennholm@ukaea.uk

Received 3 February 2024, revised 8 May 2024
Accepted for publication 8 July 2024
Published 8 August 2024

Abstract

In 2019 the UK launched the Spherical Tokamak for Energy Production (STEP) programme to design and build a prototype electricity producing nuclear fusion power plant, aiming to start operation around 2040. The plant should lay the foundation for the development of commercial nuclear fusion power plants. The design is based on the spherical tokamak principle, which opens a route to high pressure, steady state, operation. While facilitating steady state operation, the spherical design introduces some specific plasma control challenges: (i) All plasma current during the burn phase should to be generated through non-inductive means, dominated by bootstrap current. This leads to operation at high normalised plasma pressure β_N with high plasma elongation, which in turn imposes effective active stabilisation of the vertical plasma position. (ii) The tight aspect ratio means very limited space for a central solenoid, imposing that even the current ramp up must be non-inductively generated. (iii) The compact design leads to extreme heat loads on plasma facing components. A double null design has been chosen to spread this load, putting strict demands on the control of the unstable vertical plasma position. (iv) The heat pulses associated with unmitigated ELMs are unlikely to be acceptable imposing ELM free operation or active ELM control. (v) To reduce and spread heat loads, core and divertor radiation and momentum loss has to be controlled, aiming to operate with simultaneously detached upper and lower divertors. (vi) High pressure operation is likely to require active resistive wall mode (RWM) stabilisation. (vii) The conductivity distribution in structures near the plasma must be carefully selected to reduce the growth rates for the vertical instability and the RWM without damping the penetration of the of magnetic fields from active control coils too much. This article describes the initial work carried out to develop a STEP plasma control system.

* Author to whom any correspondence should be addressed.

Keywords: spherical tokamak, fusion power plant, plasma control, bootstrap current, detachment, double null, vertical stabilisation

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

1. Introduction

In a major drive to further the development of nuclear fusion as a viable energy source for the future, the UK is putting significant resources into the 'Spherical Tokamak for Energy Production' (STEP) project [1–3]. This project aims at building a prototype fusion power plant by ∼2040. The prototype should demonstrate electricity production and it should pave the way for the building of commercial nuclear fusion power plants. The project is currently in the conceptual design phase and it will progress towards detailed design over the next years, aiming to start construction in the early 2030s.

A spherical tokamak design has been chosen as such a design shows the promise of operating at higher normalised plasma pressure:

$$\beta_{\mathrm{N}} = 2\mu_0 \frac{\langle p \rangle a}{B_{\mathrm{t}} I_{\mathrm{p}}}, \tag{1}$$

where a is the minor radius, B_{t} is the toroidal field, I_{p} is the plasma current and $\langle p \rangle$ is the volume-averaged pressure [4]. Operation at high beta has the double advantage of allowing the design of a more compact machine and of operating such a machine with a high 'bootstrap current' fraction, where the bootstrap current is an autogenerated plasma current, which depends on the gradient of the plasma pressure [5].

Having a large fraction of the plasma current driven by this bootstrap mechanism, means that steady state operation becomes feasible without the need for excessive external input power to drive the plasma current. High pressure operation is facilitated by having a high plasma elongation and the spherical design allows operation at higher elongation than more conventional tokamak designs. Elongated plasmas—plasmas which are taller than they are wide—are vertically unstable, meaning that the plasma vertical position will accelerate exponentially either up or down without active stabilisation. The higher the elongation, the higher the instability growth rate. The maximum elongation which can be stabilised is a function of the plasma geometry and the conductivity of the surrounding structures. This maximum increases with decreasing plasma aspect ratio—highlighting the advantage of the low aspect ratio spherical design.

The advantage of the spherical design comes at the price of more severe plasma control challenges [6], the stabilisation of the vertical position being just one of them. The small aspect-ratio means that there is little space available for the central column. This can be seen in figure 1, which shows a sketch of the present layout of the machine including the poloidal contour of the plasma facing components, poloidal field coils and a plasma surface contour. The central column has to accommodate the toroidal field coils and the central solenoid. Note that this design is evolving, and the figure is only for illustration. Given the lack of space, the central solenoid will be of very limited size and the aim is to only use this solenoid for plasma formation (breakdown) and the early phase of plasma current ramp up. After this early phase the plasma current will be driven by purely non inductive means, initially dominated by external current drive, transiting to bootstrap current dominated plasmas in the steady state phase.

In the energy producing phase the alpha particles generated by the fusion reaction will be the dominant heating source. Thus, the plasma will be a highly self-organising system, governed by bootstrap current and alpha particle heating. The plasma control system must be able to guide and stabilise this system using the limited external actuators. Four fifth of the power produced in the fusion reaction will leave the plasma in the form of 14 MeV neutrons and these will be absorbed in the breeding blankets surrounding the plasma, producing the tritium required for continuous operation [7]. The remaining fifth of the fusion power plus the external power injected into the plasma, will leave the plasma impacting the plasma facing components. Without further measures this power will exit the plasma mainly as charged particles, which will be guided by the magnetic configuration towards the divertor strike-points where these particles would impact narrow annular surface regions. The power density in these regions would exceed material limits and hence the following measures are planned:

(i) Significant power should be leaving the core plasma as electromagnetic radiation. This can be achieved by injecting appropriate impurities into the plasma core. The aim would be to provoke radiation from the outer part of the core, just inside the last closed plasma surface, while limiting the reduction in fusion reactivity caused by radiation from the plasma centre. Good control will be required to strike the optimal balance between fusion performance and heat load distribution.

(ii) The remaining power will be leaving the plasma as charged particles which will be guided into the divertor regions. To distribute the power as evenly as possible over the divertors, the plasma will be operated in the so-called 'detached' regime, where the plasma energy is dissipated through radiation and momentum and particle loss with only low energy particles impacting the strike-points. This radiation is controlled by the injection of fuelling and impurity gas into the divertor regions [8, 9].

(iii) Finally, the design foresees to operate in a 'double null configuration', with an upper and a lower divertor. This

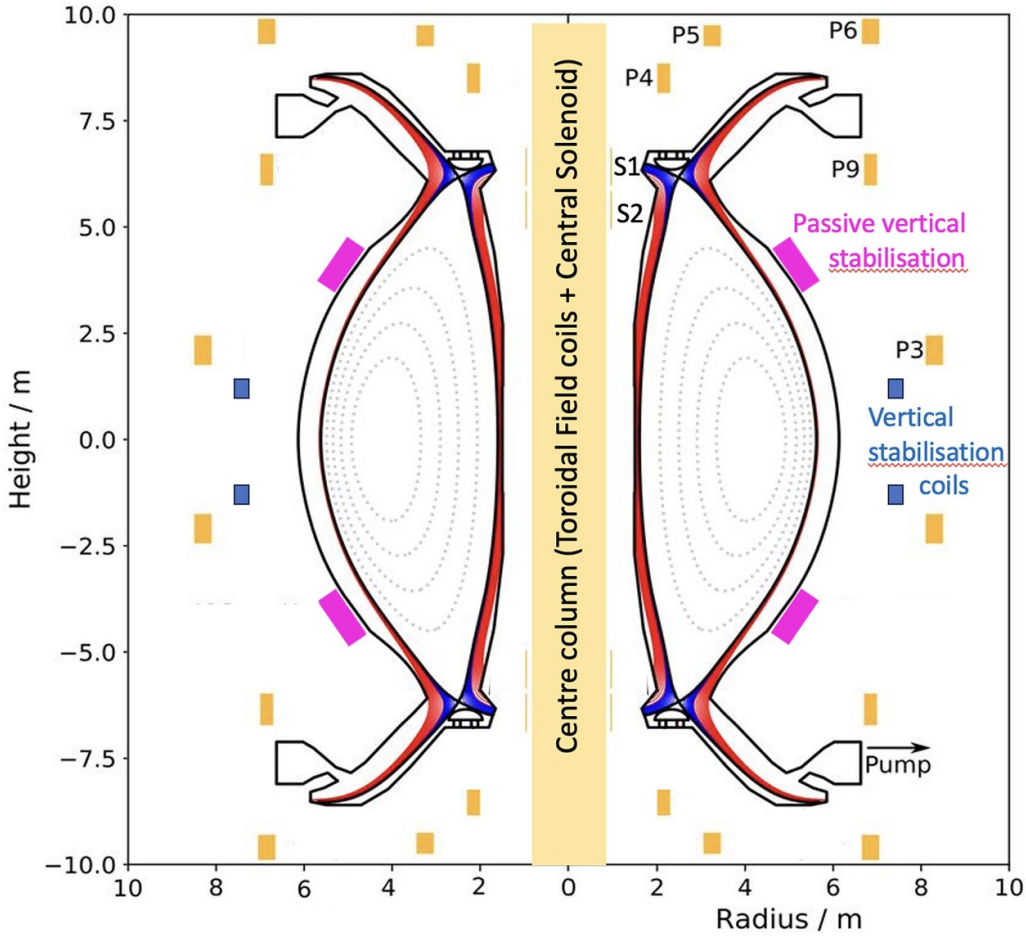

Figure 1. Sketch based on the present STEP prototype design. Note that this design is evolving, and the figure is for illustration only. The figure includes PF coils, passive structures, plasma facing component contour and an example equilibrium.

will approximately double the surface area available to handle heat and particle exhaust. In addition, double null operations has the promise of further reducing the power to the more critical inner strike-points, by isolating these strike-point from most of the power conducted across the last closed flux surface. To achieve maximum benefit from this configuration, accurate control of the vertical plasma position is required [10–14].

To achieve sufficient confinement, the STEP prototype should operate in the high confinement H-mode. This mode is normally associated with periodic energy pulses leaving the plasma core through the Edge Localised Mode (ELM) [15, 16]. Without mitigation, the energy density in such ELMs is predicted to exceed the divertor capacity, with energetic particles 'burning through' the detachment region and reaching the material surfaces. H-mode operation without ELMs has been demonstrated and such operation.would be highly desirable [17]. As the experimental foundations for such operation does not allow confident extrapolation to the STEP prototype plasmas, specific toroidally non-symmetric (3-D) 'Resonant Magnetic Perturbation' (RMP) coils aimed at mitigating ELMs are included in the design [18, 19]. The details of the work on ELM mitigation is beyond the scope of this paper.

The plan to operate at high plasma pressure, means that 'Resistive Wall Modes' (RWMs) may be unstable. To assure that such modes can be stabilised, a fast closed loop control system based on a set of 3D RWM coils is included in the STEP prototype design [20].

In addition to these control tasks, which are specific to spherical tokamaks, the plasma control system also has to handle the more traditional tasks of controlling the plasma shape and position using the set of resistive and superconducting poloidal field coils shown in figure 1.

Within the vacuum vessel, almost all components will be electrically conducting to a larger or smaller degree. The distribution of conductivity within these structures strongly affects the plasma stability and the ability to control the plasma using magnetic coils. This conductivity is both beneficial and detrimental. In fact, both the Vertical instability and the RWM instability are stabilised by currents induced in the conducting structures and neither would be stabilisable without the stabilising effect of the conducting structures. On the other hand, these structures hamper the penetration of control fields from the active control coils to the plasma and they also damp the magnetic field variation reaching magnetic sensors which may

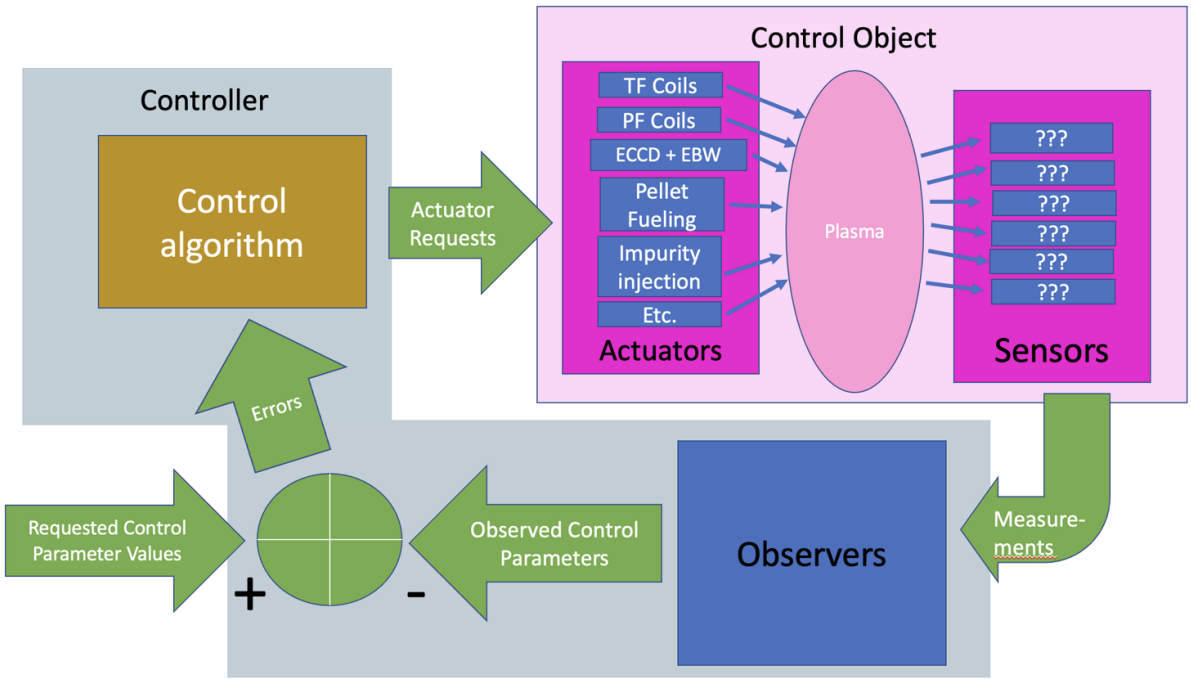

Figure 2. General plasma control loop.

be used to generate the measurements required by the plasma position and RWM control. Furthermore, too high conductivity allowing significant up-down symmetric toroidal currents to be induced in these structures will prevent plasma initiation. As a consequence, careful design of the conductivity distribution in the material surrounding the plasma is required. The freedom to design the conductivity distribution is hampered by the lack of isolating materials which can survive the neutron bombardment expected in these areas.

The high neutron rate in a fusion reactor during the burn phase puts significant restrictions on the diagnostic systems, which can be used for measuring plasma behaviour. It is unlikely that it will be possible to measure all the plasma parameters which would be valuable to enable good plasma control. To deal with this challenge, real time dynamic plasma models, constrained and continuously updated to agree with the available measurements, will be used to generate the required plasma 'measurements' to be used by the control systems. These models will be improved and tuned in an early low neutron rate operational phase where significantly more measurements will be available and the transition to high neutron rate operation will only proceed when confidence has been gained that the machine can be operated with the reduced set of 'neutron tolerant' diagnostics.

While the choice of diagnostics for the STEP prototype is still at an early stage, the actuators which will be available to the plasma control system have received more attention. These actuators fall in 3 main categories: (1) Superconducting and resistive magnetic field coils (poloidal field coils, 3D coils). (2) Heating and current drive systems. Here Electron Cyclotron heating and Current Drive (ECCD) and Electron Bernstein Wave Current Drive (EBW) has been chosen [21, 22]. (3) Pellet and gas matter injection (Fuelling and Impurities).

Though the various control tasks are discussed separately, it is clear that all the controllers will interact with each other and the plasma control system is a large Multi-Input Multi-Output (MIMO) system as illustrated in figure 2.

It is tempting to consider that having designed a control system that is capable of initiating the plasma, ramping it to burn, controlling the burn and then terminating the plasma again, the job is done. However, the most challenging task of the control system is probably associated with situations when things do not go as smoothly as desired. We discuss this in terms of 'exception handling'. When an exception is detected the plasma control system have to include logic to decide how to proceed.

Similar design work to that described below is being carried out both for ITER [23] and DEMO [24] though significant differences exist, both due to the specific characteristic of STEPs spherical tokamak approach and due to the aim to operate in steady state. This being said, it is clear that collaboration on plasma control work between the different reactors under design is highly desirable.

The remainder of this paper is organised as follows: section 2 discusses the use of real time models for control in general terms; section 3 deals with exceptions handling while section 4 goes into more details of the following specific control tasks and their interaction: Plasma profile control, Plasma Vertical Stabilisation, RWM control and Detachment control. Section 5 then draws some conclusions.

2. Use of models in control and estimation

Significant challenges will be encountered when designing the measurement systems for the STEP prototype. The magnetic

sensors which are traditionally used by the plasma shape and position control systems will be situated behind the neutron absorbing and tritium breeding blanket. Convolution effects due to the blanket and due to the larger distance between plasma and magnetic sensors will reduce both temporal and spatial resolution. Some traditional measurement systems may be completely incompatible with the high neutron dose, while some systems will be highly restricted. In addition, the measurement environment is likely to be very noisy, both due to neutron irradiation of sensors and due to vibrations. To deal with the limitations of available sensors and measurement noise, it is proposed to use model-based state observers to determine the state variable evolution of the system in real time. In this scheme, models of the control object will compute the expected system state vector $\boldsymbol{S}_{\mathrm{MOD}}(t)$ as a function of time t, based on the history of actuator requests represented by the vector $\boldsymbol{I}(t)$:

$$\boldsymbol{S}_{\mathrm{MOD}}(t) = F_s(\boldsymbol{I}(t), \boldsymbol{I}(t-\Delta t), \ldots, I(t-m\Delta t), \boldsymbol{S}_{\mathrm{Mod}}(t), \ldots, \boldsymbol{S}_{\mathrm{Mod}}(t-m\Delta t), t, \Delta t, m, \boldsymbol{P}) \tag{2}$$

where F_s is a function which represents how the model derives the expected state evolution as a function of the actuator and state history. Here $\boldsymbol{P}$ is a vector of internal model parameters, Δt is the time step used by the model and m is the number time steps considered by the model.

The model also gives a prediction for the vector of diagnostic measurements as a function of time (synthetic diagnostic):

$$\mathbf{D}_{\mathrm{MOD}}(t) = F_D(\mathbf{I}(t), \mathbf{I}(t-\Delta t), \ldots, \mathbf{I}(t-m\Delta t), \mathbf{S}_{\mathrm{Mod}}(t), \ldots, \mathbf{S}_{\mathrm{Mod}}(t-m\Delta t), t, \Delta t, m, \boldsymbol{P}) \tag{3}$$

where F_D is a function which represents how the model derives the expected measurement vector.

This computed sensor response will be compared to the actual measurements from the sensors

$\mathbf{D}_{\mathrm{MEAS}}(t)$ giving rise to the model error:

$$\mathbf{E}(t, \boldsymbol{P}) = \mathbf{D}_{\mathrm{MOD}}(t, \boldsymbol{P}) - \mathbf{D}_{\mathrm{MEAS}}(t). \tag{4}$$

From this, a cost $J(\boldsymbol{P})$ can be derived, showing how well the model matches reality as a function of $\boldsymbol{P}$

$$J(\boldsymbol{P}) = F_C(\boldsymbol{E}(t, \boldsymbol{P}), \boldsymbol{E}(t-\Delta t, \boldsymbol{P}) \ldots, \mathbf{E}(t-n\Delta t, \boldsymbol{P})) \tag{5}$$

where F_C is a function giving the cost as a function of $\boldsymbol{P}$ and $n\Delta t$ is the time horizon for the cost calculation.

Now the model can be optimised by finding the value $\boldsymbol{P}_{\mathbf{opt}}$ of $\boldsymbol{P}$ which minimises J:

$$\{\boldsymbol{P}_{\mathbf{opt}} \in \Omega_{\boldsymbol{P}} | \forall \boldsymbol{P} \in \Omega_{\boldsymbol{P}}, \ |J(\boldsymbol{P}_{\mathbf{opt}})| \leqslant |J(\boldsymbol{P})|\} \tag{6}$$

where $\Omega_{\boldsymbol{P}}$ is the subspace of values of $\boldsymbol{P}$ over which the minimisation is carried out. It is obvious that, the larger $\Omega_{\boldsymbol{P}}$ is, the longer this optimisation will take and hence $\Omega_{\boldsymbol{P}}$ has to be chosen carefully. It is likely that $\Omega_{\boldsymbol{P}}$ will form a neighbourhood around the $\boldsymbol{P}_{\mathbf{opt}}$ for the previous timestep and hence $\Omega_{\boldsymbol{P}}$ will evolve from timestep to timestep. Standard methods exist for such cost function minimisation as discussed further in [25].

This process assures that the model is in good agreement with reality. The model with the optimised parameter set now give a best possible real time estimate of the state variables of interest

$$\mathbf{S}_{\mathrm{MOD\_opt}}(t) = F_s(\boldsymbol{I}(t), \boldsymbol{I}(t-\Delta t), \ldots, \boldsymbol{I}(t-m\Delta t), t, \Delta t, m, \boldsymbol{P}_{\mathbf{opt}}). \tag{7}$$

In this way the model acts as a state observer which is constrained by the available measurements. The observed state variables can then be used in a traditional control scheme by computing the control error:

$$\mathbf{E}_{\mathrm{CONT}}(t) = \mathbf{S}_{\mathrm{REF}}(t) - \mathbf{S}_{\mathrm{MOD\_opt}}(t) \tag{8}$$

where $\mathbf{S}_{\mathrm{REF}}(t)$ is the required state evolution. The controller will use the error including previous time points to derive the actuator request:

$$\mathbf{I}(t+\Delta t) = F_{\mathrm{CONT}}(\mathbf{E}_{\mathrm{CONT}}(t), \mathbf{E}_{\mathrm{CONT}}(t-\Delta t), \ldots, \mathbf{E}_{\mathrm{CONT}}(t-n\Delta t), t, \Delta t, n) \tag{9}$$

where n is the number of previous time samples considered by the controller and F_{CONT} is the control function.

More advanced control options become available when the models can run 'faster than real time'. In this case the model is used to derive the future trajectories of the state variables $\mathbf{S}_{\mathrm{MOD\_opt}}(t+\Delta t, \mathbf{I})$, $\mathbf{S}_{\mathrm{MOD\_opt}}(t+2\Delta t, \mathbf{I}) \ldots$ $\mathbf{S}_{\mathrm{MOD\_opt}}(t+k\Delta t, \mathbf{I})$ as a function of the future actuator evolution $\mathbf{I} = \mathbf{I}(t+\Delta t)$, $I(t+2\Delta t) \ldots I(t+k\Delta t)$. Here k represents the number of time steps that the model predicts. The resulting state vector trajectories are then compared to the requested trajectories $\mathbf{S}_{\mathrm{REF}}(t+\Delta t), \mathbf{S}_{\mathrm{REF}}(t+2\Delta t) \ldots \mathbf{S}_{\mathrm{REF}}(t+k\Delta t)$ to generate the trajectory error:

$$\begin{aligned}
&\mathbf{E}_{TRA}(t+\Delta t, \mathbf{I}) = \mathbf{S}_{\mathrm{REF}}(t+\Delta t) - \mathbf{S}_{\mathrm{MOD_opt}}(t+\Delta t, \mathbf{I}) \\
&\mathbf{E}_{TRA}(t+2\Delta t, \mathbf{I}) = \mathbf{S}_{\mathrm{REF}}(t+2\Delta t) - \mathbf{S}_{\mathrm{MOD_opt}}(t+2\Delta t, \mathbf{I}) \\
&\vdots \\
&\mathbf{E}_{TRA}(t+k\Delta t, \mathbf{I}) = \mathbf{S}_{\mathrm{REF}}(t+k\Delta t) - \mathbf{S}_{\mathrm{MOD_opt}}(t+k\Delta t, \mathbf{I}).
\end{aligned} \tag{10}$$

From this we can generate a cost showing how well the model predicted state trajectory matches the required future trajectory:

$$J_{TRA}(\boldsymbol{I}) = F_{TRA}(\boldsymbol{E}_{TRA}(t+\Delta t, \mathbf{I}), E_{TRA}(t+2\Delta t, \mathbf{I}) \ldots, E_{TRA}(t+k\Delta t, \mathbf{I})). \tag{11}$$

Now the trajectory can be optimised by finding the values $\boldsymbol{I}_{\mathbf{opt}}$ of $\boldsymbol{I}$ which minimises J_{TRA} within the space allowed for the trajectory search Ω_I

$$\{\boldsymbol{I}_{\mathbf{opt}} \in \Omega_{\boldsymbol{I}} | \forall \boldsymbol{I} \in \Omega_{\boldsymbol{I}}, \ |J(\boldsymbol{I}_{\mathbf{opt}})| \leqslant |J(\boldsymbol{I})|\}. \tag{12}$$

Finally, $\boldsymbol{I}_{\mathbf{opt}}(t+\Delta t), \boldsymbol{I}_{\mathbf{opt}}(t+\Delta t)$ …. can be applied to the actuator. This actuator trajectory optimisation could be applied

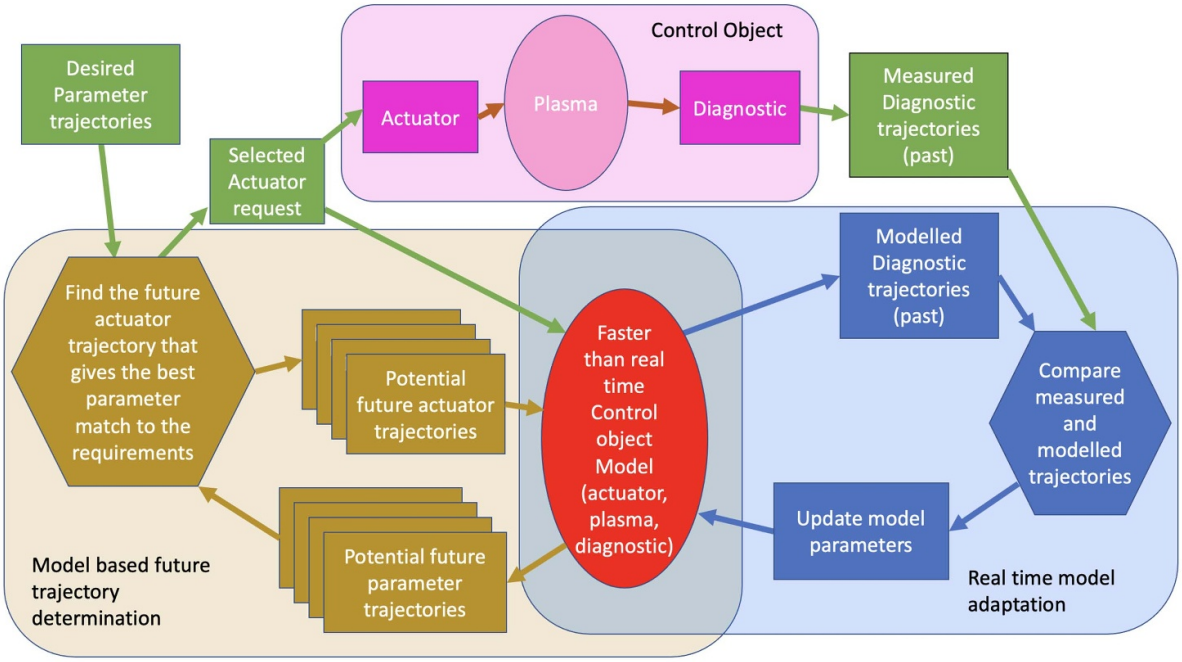

Figure 3. Model predictive control—scheme.

at each time step or less frequently depending on the computational demand. The larger a future parameter space that is searched and the further into the future the match is sought, the more computationally demanding this process gets. Taking this process a step further, a similar optimisation procedure can be implemented to search for an optimal trajectory in parameter space, in order to achieve a specific goal. In this case constraints can be imposed within plasma and actuator parameter space and a cost function can be calculated as a function of the trajectory. This cost is then minimimised while respecting the constraints, to find the best way to reach a specific goal. Such optimisation can be done prior to the start of a plasma or, computational power allowing, continuously throughout a plasma discharge combined with a continuous updating of the goals. An example of the use of such trajectory optimisation is discussed in section 4.1.

Figure 3 illustrates the principle discussed above. Here the red oval represents the 'faster than real time' model. The light blue area represents the continuous model adaptation comparing measured and simulated sensor data. The light brown area represents the selection of future actuator trajectories to achieve the best possible match between predicted and required state variable evolutions. Such control schemes are likely to be essential given the slow response of the system to actuator variations. Figure 4 shows a tentative illustration of such a scheme for one input and one controlled variable. For simplicity it is assumed that the controlled variable is measured directly. The top traces show the actuator value evolutions while the bottom traces show the controlled variable evolutions. Up until the 'now' time, the magenta trace shows the value of the controlled variable as computed by the model, while the continuous orange trace shows the corresponding measurement. The jumps in the magenta line illustrate the model being adapted at each time step to assure match between simulated and measured values. After the 'now' time, five options for potential actuator trajectories are shown in the top traces and the corresponding evolutions of the controlled parameter as predicted by the model are shown in the bottom traces. Based on comparing these traces with the desired parameter evolution shown as a black dashed curve, the controller would choose the red traces and send the corresponding requests to the actuator. This process could be repeated at each time step or, if computational limitations does not allow this, the predictive computation could be repeated less frequently than the model adaptation. The prediction time horizon has to be chosen carefully taking computational time and model confidence into account.

The predictive capability of this process can be taken a step further by searching for optimal future state evolutions to achieve some higher level goals. In this way parameter and actuator limitations can be taken into account and a range of potential trajectories can be explored to find an optimal solution. This optimisation can be repeated as frequently as computational power allows. Such a process will also allow real time decision making on actuator allocation in cases where individual actuators may serve more than one purpose [26].

For a high dimensional system with many actuators, sensors and states such a controller will rapidly require significant computational resources and it is therefore important to develop models that are not too computationally demanding. This model reduction has to be weighed up against the precision required from the output from the model. The

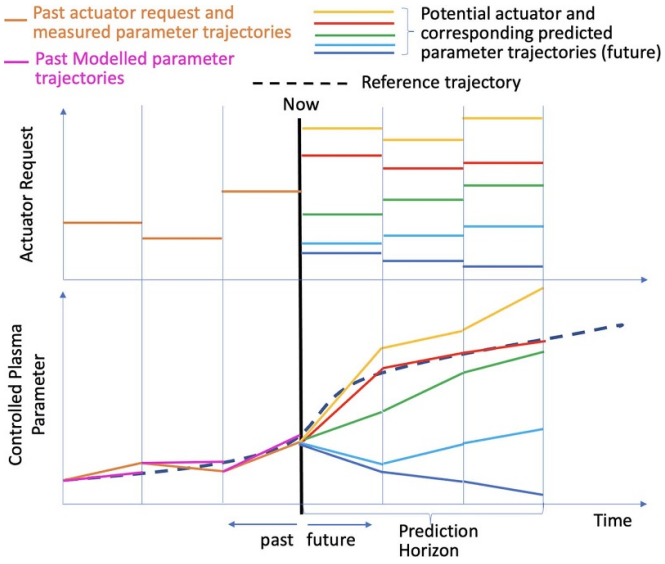

Figure 4. Model predictive control—sketch of actuator and parameter trajectories for a single actuator and 1 parameter.

better the measurement coverage, the less stringent the model requirements. As the number of measurements is reduced and the measurement noise increases, the required model fidelity increases. In the design phase the best trade-off between model and measurement requirements should be found. Given the uncertainty associated with modelling the plasma, the plan is to equip the machine with as many diagnostics as possible for the low radiation phase. In the design phase it is important to assure that sufficiently reliable models can be developed by learning from operation in this phase to be confident that these models will allow robust operation with the limited sensor set that will be available in the high activation phase.

The operational plan for the STEP prototype involves an initial low neutron rate phase with high diagnostic coverage, where all machine systems can be brought into operation. Following this initial phase, the neutron radiation will be increased stepwise until full fusion power operation is reached. Before each increase in neutron radiation, it is important to demonstrate reliable operation with the reduced sensor set which will be available after the increase.

3. Exception handling

While the above section deals with nominal operation, an essential and more challenging task for the plasma control system is to respond correctly when things do not go exactly as desired. We term this 'exception handling'. Exceptions that the plasma control system must deal with cover a very wide range from rather benign events which hardly affect the operation, to severe events, which eventually may lead to machine damage. A long list of exceptions has to be handled and the required action depends on the type and severity of the exception. Exceptions that should be handled include both failure of equipment and undesirable plasma behaviour. Examples of equipment issues include failures or limitations of actuators, sensors, communications networks, software etc. Other exceptions could be excessive heating of plasma facing components, growth of undesirable magneto hydrodynamic (MHD) modes or impurity accumulation in the plasma core. A hierarchy of exceptions can be defined. The most severe plasma event is a high current plasma disruption and one of the most important goals of the exception handling system is to prevent disruptions. Failing in this task, the system has to minimise the disruption impact [27]. An exception will be characterised by a parameter or a combination of parameters exceeding some threshold. Many different thresholds will exist. In an abstract way, each exception will be characterised by a Boolean variable Ex_n where $n < N$ is an integer number and N is the number of different exceptions which need to be handled. Whether Ex_n is true or false (i.e whether exception n is ongoing or not) depends on the state of the system represented by the state vector $\boldsymbol{P}$ and a vector of thresholds T as expressed through the real function F_{Ex_n}

$$Ex_n = F_{Ex_n}(\boldsymbol{P}, \boldsymbol{T}) > 0. \tag{13}$$

The response required will depend on the value of Ex_n for all $n < N$. Thus a number M of responses Res_m with $m < M$ will be defined and a Boolean variable B_{Res_n} will be calculated for each of these responses where the value of B_{Res_m} (TRUE or FALSE) determines whether Res_m is active or not

$$B_{Res_m} = FB_{Res_m}(Ex_1, Ex_2, \ldots, Ex_N). \tag{14}$$

The specifics of each response will range from simple responses using a single actuator to much more complex responses. The speed of the response will depend on the risk associated with the exception, the time required to detect the event and the responsiveness of the actuators which will be used. In the case of high urgency, simple detection and fast actuators have to be prioritised and the exception handling algorithms cannot be allowed to delay the response significantly. When a slower response is adequate, it is likely that real time calculation of trajectories in plasma and actuator parameter space will be used to handle each exception in the best possible way.

Figure 5 attempts to illustrate this multidimensional control and exception handling task by looking at a projection into a 2-dimensional space. Here the two axes represent 2 parameters/states, which need to be controlled to achieve the desired plasma behaviour. This is just an illustration for a 2-dimensional subset of a high dimensional space, which includes all the parameters that are important for plasma stability and control. These parameters are not limited to plasma parameters, but will also include actuator and sensor availability parameters, temperatures of plasma facing components etc. For successful operation the controllers need to assure that the controlled parameters follow a complex trajectory in this multidimensional space avoiding unstable or undesirable regions. In figure 5 the black trajectories represent the nominal trajectories which it is desired that the plasma follows. The ramp up trajectory goes from plasma initiation (represented by a triangle) to the nominal operations point (represented by a circle), where the control system should maintain it for

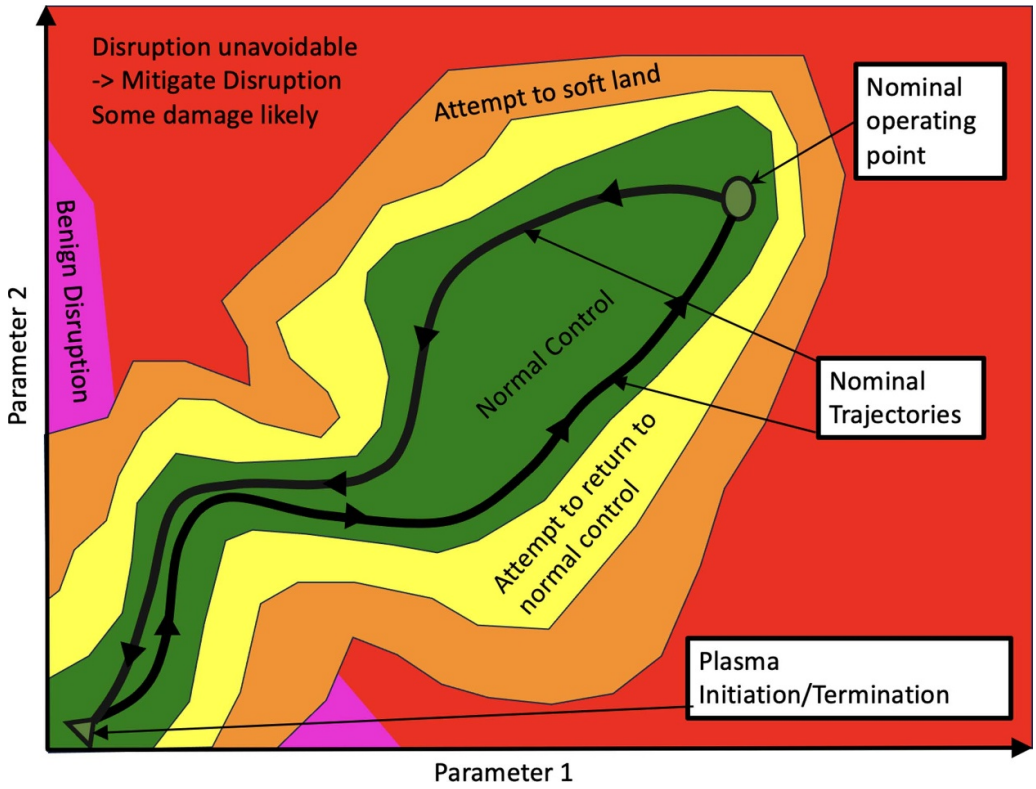

Figure 5. Conceptual 2D illustration of parameter space and the response of the plasma control system to exceptions.

the duration of the steady state phase [28]. The second trajectory shows the desired termination trajectory, which is likely to differ significantly from the ramp up trajectory. Minor disturbances and noise will mean that these trajectories are not followed precisely but, as long as the deviation is not too large, this is still considered normal operation. The part of operational space, which is considered normal is represented by the green area in figure 5. While operating in this area the controller will continuously update the models and tune the trajectories as described in section two.

If the system is disturbed sufficiently to leave the green area an exception handling response as defined above, is required. If the exception or combination of exceptions is minor—and recoverable—the system will be considered to have entered the yellow area. A range of different responses will be associated with this yellow area, but common to them all are that they aim to bring the plasma back to 'normal'. In this case the control system determines a safe path back into the normal space and then act to follow this path. Such minor exceptions could be loss of sensors or actuators, where sufficient redundancy exists. An example of such loss of redundancy would be a reduction in available ECCD/EBW power. Though virtually transparent to the control system, a reduction of the available redundancy still should be considered as an exception as this may mean that a further loss of sensor or actuator capability can have a more significant impact and the exception handling system has to determine whether returning to normal operation is the correct action.

If the combination of exceptions is such that a return to normal operation is not feasible, or unsafe, the system will be considered to have entered the orange area in the figure. In this area, it is expected that the plasma can be terminated safely. Similarly to the yellow area, a range of different responses, as defined above, are associated with the orange area. Common to these responses is the aim of terminating the plasma safely and the responses will seek to determine the best trajectory in parameter space to assure such a termination.

If an exception is sufficiently severe, this is illustrated by the red area. For this area, a disruption is deemed to be unavoidable. In this case, the response of the exception handling system aims to mitigate the disruption as best possible. The actions undertaken are likely to include the following:

(i) reduce the plasma current and plasma energy as fast as possible,
(ii) fire shattered pellets into the plasma to radiate the plasma energy over the largest possible area,
(iii) assure that the remaining plasma energy is deposited on less sensitive components like 'sacrificial' limiters.

The pink area in figure 5 represents plasmas which are going to disrupt, but where the plasma current and energy are small enough for the impact of such disruptions to be benign. If the exception handling system manages to bring an unhealthy plasma into this area before it disrupts, this is considered as a successful termination.

From this discussion it is evident that the severity of the combination of exceptions can escalate from yellow, via orange to red, if successive responses prove unsuccessful.

A supervisory protection system will exist independently of the exception handling component of the plasma control system. This supervisory system will be able to initiate emergency action, overruling the plasma control system if needed. One of the aims of the exception handling system, within the plasma control system, is to avoid triggering emergency action from the supervisory protection system.

A few examples of exceptions and responses associated with the control challenges discussed in section 4 have been added to the specific sections. It should, however, be noted that many exceptions and their handling will involve most or all of the individual controllers.

4. Specific control challenges

This section describes a selection of the work which has been carried out to further explore specific control issues which have been deemed particularly demanding in a fusion reactor based on the spherical tokamak principle. One of the main tasks which a plasma control system has to address in any tokamak is assuring that the plasma shape and position is as required. This is a well established control task where existing solutions are likely to be directly applicable to the STEP design. A couple of exceptions to this should be mentioned. The control of the unstable vertical position in the STEP double null plasma will be challenging. More generically, the shape and position control system will suffer from limited diagnostic available affecting all the plasma controllers [29]. Given these considerations less effort has been dedicated to the general plasma position and shape controller and the main effort has been dedicated to working on the vertical position control part of this controller. As a consequence, with the exception of vertical position control, the plasma shape and position control system is not discussed in further details below. Significant efforts will be dedicated to this over coming years, once other controller design tasks have reached a similar level of maturity. The following sections describes the main areas which the STEP plasma control team has focused on over recent years.

4.1. Profile control of self-organising system

It is highly desirable that the STEP prototype power plant can be operated in steady state. This will significantly reduce the thermal cycling of plant components and it will mean that the energy used during plasma initiation, ramp up and termination becomes negligible in comparison with the energy produced during the extended steady state phase. The main difficulty associated with operating a tokamak in steady state, is to maintain the required plasma current. In traditional tokamak operation, the plasma current is generated through transformer action, by continuously varying the current in a central solenoid. Such operation is inherently pulsed and hence steady state operation requires alternative ways of driving the plasma current. Plasma current can be driven by various heating and current drive systems. For STEP, a combination of electron cyclotron current drive (ECCD) and electron Bernstein wave current drive (EBW) has been chosen [21]. Driving the full required plasma current using these external means would demand very significant electrical power, significantly impacting the ability to produce net electricity. To reduce the required ECCD/EBW power, the aim is for most of the STEP current to be generated through the 'bootstrap' mechanism. The spherical design allows operation at higher plasma elongation, higher normalised plasma pressure and pressure gradient and thereby with higher bootstrap fraction. The aim is to design for most ($>80\%$) of the plasma current to be generated through this mechanism during the burn phase. During this phase, the resulting plasma will constitute a highly self-organising system, with the plasma being mainly heated by alpha particles created by the fusion reaction. The alpha heating in turn dominates the pressure profiles and through this, the plasma current profile. The current profile itself strongly affect the plasma confinement and plasma instabilities and thereby the generated fusion power and alpha particle heating. It is the task of the plasma control system to break into this system via the limited available actuators to ensure that the plasma current and pressure profiles remain stable, while optimising the fusion power and bootstrap current generation.

Due to the minimal space available for the central solenoid almost all the plasma current, even throughout the plasma current ramp up, has to be generated by non-inductive means. During the ramp-up the alpha power will be moderate and hence the plasma current has to be ramped up through the use of the ECCD/EBW. To achieve high ECCD/EBW current drive efficiency during the ramp up phase operation at low plasma density is required in contrast to the burn phase, where high density is required to produce high fusion and alpha power.

Finding the best trajectory for the plasma parameters to move from plasma initiation to a steady state burning plasma is a challenging multidimensional optimisation problem. The predictive plasma simulation tool RAPTOR has been used to find such optimal trajectories for different optimisation criteria [25, 30, 31]. To do this, the prediction horizon has been chosen to be the entire time from plasma initiation to steady state burn with the code minimising a cost function within a range of constraints. The non-linear RAPTOR code is a simplified transport code allowing it to run fast enough to execute complex optimisation tasks within a manageable time. In these optimisations tasks Bohm-GyroBohm transport scaling has been assumed. Note that robust controllers and the use of model adaptation as described in section 2 should be able to achieve good plasma control even for plasmas with varying transport properties. Once RAPTOR has found promising parameter trajectories, these trajectories are verified with higher fidelity,

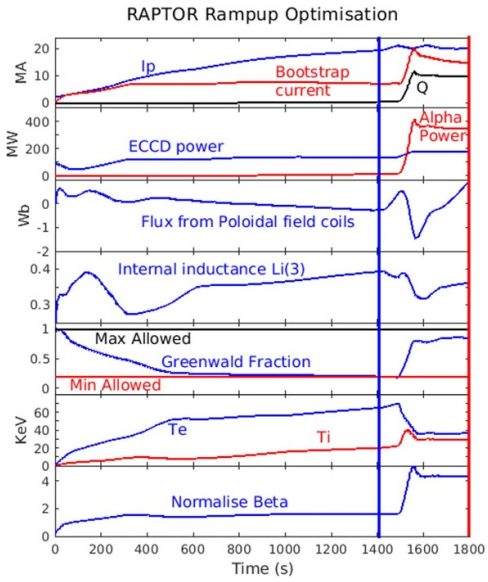

Figure 6. RAPTOR optimised ramp-up trajectories. Blue and red vertical lines show the time points corresponding to profiles shown in figure 7.

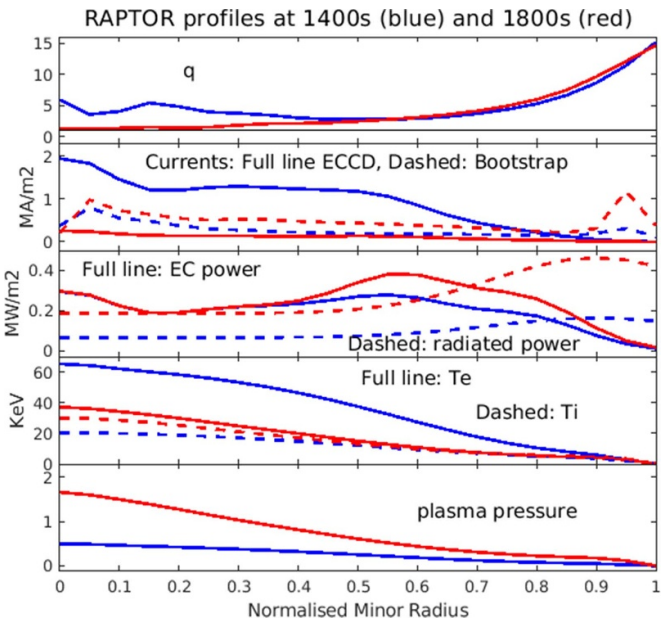

Figure 7. Plasma profiles before and after density increase, showing the move from ECCD dominated to bootstrap dominated plasma.

more computationally demanding, codes such as JINTRAC [32, 33]. Figure 6 shows a potential ramp up scenario as found using the RAPTOR trajectory optimisation capabilities. This is done by minimising the following cost function:

$$J_{\text{tot}} = J_{\varphi\,\text{edge}} + J_{\text{Li(3)}} \tag{15}$$

where

$$J_{\varphi\,\text{edge}} = W_{\varphi\,\text{edge}} \int_0^{t_0} [\varphi_{\text{edge}}]^2 \mathrm{d}t \tag{16}$$

$$J_{\text{Li(3)}} = W_{\text{Li(3)}} \int_0^{t_0} [\text{Li}\,(3) - 0.3]^2 \mathrm{d}t. \tag{17}$$

Here φ_{edge} is the edge flux, which reflects the use of the central solenoid and Li*(3)* is the internal inductance which should be kept near 0.3. $W_{\varphi\,\text{edge}}$ and $W_{\text{Li(3)}}$ are the weights applied to set the relative importance of maintaining the edge flux and the internal inductance near the desired values. t_0 is the time horizon for the optimisation. The optimisation is done while imposing the following constraints:

$$C_{F_{\text{GWU}}} = W_{F_{\text{GF}}} \int_0^{t_0} (\max\,[0, F_{\text{GW}} - F_{\text{GWmax}}])^2 \mathrm{d}t \tag{18}$$

$$C_{F_{\text{GWL}}} = W_{F_{\text{GF}}} \int_0^{t_0} (\max\,[0, F_{\text{GWmin}} - F_{\text{GW}}])^2 \mathrm{d}t \tag{19}$$

$$C_q = W_q \int_0^{t_0} \left[\int_0^{\rho_0} \max\left[0, \frac{1}{q} - 1\right] \mathrm{d}\rho \right]^2 \mathrm{d}t \tag{20}$$

$$C_{T_e} = W_{T_e} \int_0^{t_0} \left[\int_0^{\rho_0} \max\,[0, T_e - 80\text{keV}]\,\mathrm{d}\rho \right]^2 \mathrm{d}t. \tag{21}$$

Here F_{GW} is the Greenwald density fraction $F_{\text{GW}} = n_e/n_{e\text{GW}}$, where n_e is the electron density and n_{eGW} is the Greenwald density limit [34], ρ_0 is the normalised minor radius within which the optimisation is carried out and q is the safety factor. These constraints are imposed as 'soft' constraints, meaning that the constraint costs are added to the costs of equation (15) before the cost is minimised. The weights applied $W_{F_{\text{GF}}}$, W_q and W_{T_e} are chosen to penalise exceeding the constraints heavily. In addition, a hard maximum ECCD power of 200 MW is imposed.

For the optimisation in figure 6, F_{GWmax} is set to 1, while F_{GWmin} is set to 0.2 until 1500s and then ramped up to 0.6 at 1550s, to force the density increase to move from ECCD dominated to bootstrap dominated operation. The constraint on the electron temperature $T_e < 80$keV is imposed to prevent the optimisation to search for solutions with unrealistic electron temperatures. The trajectory found consists in a plasma current ramp up at low density followed by a rapid increase in density, transitioning into the bootstrap dominated burn phase. Figure 7 shows the radial profiles of safety factor (q), ECCD current, Bootstrap Current, Radiated power, Electron and Ion temperatures and plasma pressure before and after the density increase. From this figure it can be seen that the final profiles are broad with significant bootstrap current generated near the edge of the plasma. Note that this is shown as an example of the trajectory optimisation capabilities of RAPTOR. A number of optimisations, using varying weights and constraints, have been carried out and the trajectories shown are likely to evolve as the machine design progresses.

This optimisation has only considered the non-inductive ramp up phase where the plasma is in a divertor configuration. Plasma initiation and current ramp up to 2 MA will be carried out using inductive current drive based on the limited central solenoid. At the end of this inductive phase the plasma will transition to a divertor configuration, before the ECCD power is increased for the non-inductive current ramp up phase. Only low power ECCD will be applied during the limiter phase to prevent overheating of the limiter.

The optimisation allows the Greenwald fraction to go to 1 throughout. This has been exploited at the beginning of the ramp up. This may be marginal, but it would be a simple matter to reduce the allowed Greenwald fraction below 1, if this is deemed necessary, without substantially changing the optimisation.

At 1400s the poloidal flux is seen to increase before the imposed density increase. This poloidal flux is available to the optimiser at a cost given by equation (16). To compensate for the increased plasma resistivity due to the density rise, a negative flux excursion will be required to keep the q profile within its constraints. The optimiser knows that this density increase is coming up and in anticipation it increases the poloidal flux and the plasma current to allow for the upcoming negative flux swing, while minimising the overall flux variation. From about 1550s the high fusion power and the high density leads to increased pressure and bootstrap current. This in turn allows the poloidal flux to increase again. Once the bootstrap current becomes dominant, much less current needs to be driven by ECCD. Despite this, more ECCD power is required during the burn phase due to the significantly reduced current drive efficiency at high density.

Simulations using RAPTOR have been undertaken to investigate the ability to control the alpha power (and thereby the fusion power) during the burn phase. Figure 8 shows the simulation of the response to a requested reduction in alpha power from 600 MW to 300 MW. In this case, the change in requested fusion power is not known to the controller beforehand. This would for example be the case if an external system or an operator requested a reduction in power. At the time of receiving this change in requested fusion power, the Model Predictive Controller (MPC) reoptimizes the proposed future actuator trajectories. The objective function is designed to minimise the tracking error for a target q profile and P_{alpha} power reference, minimise the ECCD and plasma current actuator usage as well as minimising edge flux consumption. This has been implemented through cost functions, similar those of equations (15)–(21). The optimisation keeps the q-profile as constant as possible by including the variation of this q-profile in the cost function. Note that in this simulation a small variation of the current in the central solenoid has been allowed inducing small transient variations in loop voltage. This simulation illustrates that a reasonably responsive fusion power controller can be developed. The simulation is done for a high flattop fusion power, which would be more relevant in a commercial fusion reactor, but which would probably be rather optimistic in the STEP prototype. Note that the power

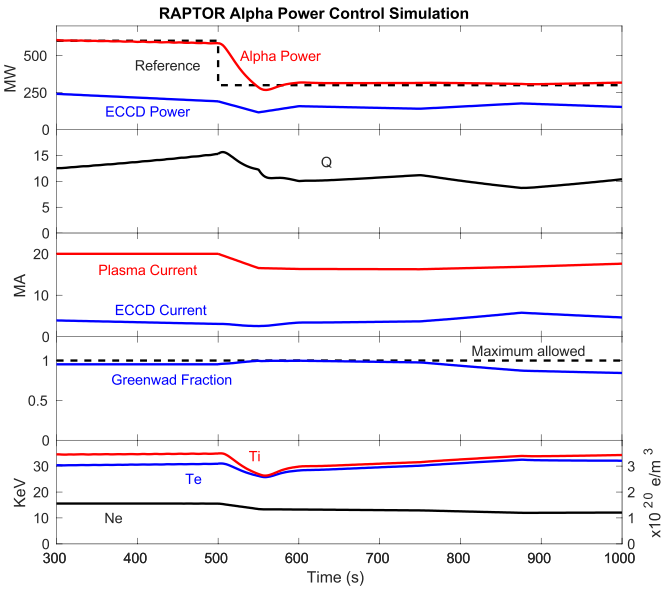

Figure 8. Burn Control: Simulated control response to a step down in requested alpha (and thereby fusion) power, resulting in lowered plasma current and density.

reduction is achieved by a reduction mainly of the plasma current and an associated reduction in density. The capital $Q = P_{\mathrm{fus}}/P_{\mathrm{aux}}$ is reduced by the stepdown but remains reasonable. Here P_{fus} is the fusion power produced by the plasma and P_{aux} is the auxiliary (heating and current drive) power injected into the plasma. Note that the initial plasma was not completely in steady state at 300s due to the very long current diffusion times, resulting in some induced counter electromagnetic field (EMF) still being present. This is the reason that the required power is slowly decreasing in this initial phase before the stepdown. More details of the RAPTOR trajectory optimisation and control simulation will be discussed in [24].

The full non-linear optimisation of the plasma current ramp up requires significant computational resource and time, rendering it unusable in real-time applications. It is, however, possible to run a linearised version of RAPTOR in real time within a model predictive controller (MPC), using a more modest predictive time horizon. It is expected that MPC of this type will form the basis for profile control on the STEP prototype.

It is expected that model-based codes, with predictive capabilities, similar to RAPTOR will be responsible for most plasma control tasks on the STEP prototype. The predictive time horizon and the update frequency have to be carefully selected to cope with measurement limitations and slow actuator response times while being able to respond sufficiently quickly to changes in requests or genuine changes in plasma conditions.

While model-based control assures some resilience against measurement limitations, this advantage is moderated by the fidelity of the models used. For this reason, it is essential to

continue to improve the modelling available for such controllers. Once the STEP prototype starts operating with the enhanced sensor set expected for the early phase, the models should be improved and updated in preparation for the less well diagnosed operation expected once the neutron radiation increases.

An example of an exception associated with the profile controller could be the loss of available ECCD power preventing adequate current profile control. In this case the response is likely to be to search for a trajectory to a stable lower plasma current operating point. The long time constants associated with the evolution of the plasma current and it is profile, are likely to allow for an effective trajectory search. Once enough ECCD power is recovered, a trajectory to return to normal operation can be sought. In this case the exception would be in the yellow region of figure 5. If sufficient ECCD power cannot be recovered the exception would enter the orange region and a trajectory allowing a safe termination should be determined. Many other potential exceptions will be considered, though their description is beyond the scope of this paper.

4.2. Magnetic control

The shape and position of the STEP plasma will be controlled via the currents in the poloidal field coils shown in figure 1. Such control is standard for all tokamaks and hence less emphasis has been given to studying this for STEP with respect to other control issues. An exception to this is the control of the plasma vertical position, which will be discussed in more details in section 4.3. The main differences between the magnetic control on STEP and conventional tokamaks are the difficulty of implementing real time sensors to determine the position and shape of the steady state plasma in the high neutron radiation environment and the use of superconducting coils for most of the poloidal field coils.

Traditionally magnetic sensing coils are used to measure the magnetic field around the plasma and based on these measurements the shape and position of the plasma is derived using an equilibrium/boundary reconstruction code [35]. The magnetic sensor coils used will measure the time derivative of the magnetic field rather than the absolute magnetic field and hence the absolute magnetic field is found through integration of these signals. For steady state operation, this integration will inevitably drift and the absolute value will become invalid. This means that alternative, absolute, measurements will be needed to compensate this drift. These alternative measurements can be slow, as their purpose will only be to compensate the integration drift. In addition to this issue related to steady state operation, the high neutron radiation means that magnetic sensors are likely to have to be situated behind significant neutron shielding, which will reduce the sensor band width due to induced currents in the shielding material. Alternative sensors are being considered though these investigations are in the early stages and will not be discussed further here.

The use of superconducting coils for control of the plasma shape and position means that this control cannot respond very rapidly as fast changes of the current in such superconducting coils will lead to significant AC losses with the risk of the coils overheating and losing their superconducting properties (quenching). Quenching such coils risks damaging or destroying the coil, so it must be avoided. This slow response of the coils means that in this early design phase, it has been considered sufficient to determine a number of static plasma equilibria throughout the plasma start up, ramp up and flat top, without considering the dynamics of the controller. The details of the dynamic control will be investigated in detail in the future. Some exceptions to this has to be mentioned. The plasma vertical position will require rapid control and to assure this, the design includes resistive, invessel, vertical stabilisation control coils (figure 1). The control of the vertical position is discussed in more detail in section 4.3. The plasma radial position is also likely to require more rapid control. If the plasma pressure changes rapidly, the plasma radial position will change significantly on a confinement timescale, due to the change in 'Shafranov Shift'. This is likely to happen when the plasma transitions between High (H-mode) and Low (L-mode) confinement. If the movement is too fast for the control coil current to compensate, significant gaps between the plasma and the plasma facing components will be required to prevent the plasma from touching these plasma facing components. As it is desirable to keep the required gap at a minimum, resistive invessel control coils are considered. These could be the same as the vertical position control coils, with the difference between upper and lower coil current, controlling the vertical position and the sum controlling the radial position. This scheme has already been proposed for DEMO [36]. For steady state operation, it is desirable to have as little current in the resistive coils as possible and the control scheme would assure that any DC/Low Frequency magnetic field components are created by the superconducting coils, with the resistive coils generating the higher frequency AC components.

The design of the plasma equilibrium evolution throughout the pulse has concentrated on the ability to initiate the plasma and ramp it to the full power configuration. Figure 9 shows some snapshots of the equilibrium throughout this phase. The first two equilibria show the limiter phase, where the current is inductively driven. The third equilibrium is early in the non-inductively driven divertor phase, while the last equilibrium show the final equilibrium during the fusion burn phase. Figure 10 shows the evolution of the currents in the poloidal field coils as a function of the plasma current corresponding to the equilibrium evolution in figure 9. Further details can be found in [10, 11].

While the rampdown phase also requires careful attention, this investigation is still in its initial phase and will not be discussed here. Suffice it to say that, as the fusion power and hence the boostrap current is reduced the current will eventually be driven purely by the ECCD/EBW systems. This increases the flexibility to control the plasma current profile and should render current ramp down easier than in purely inductively driven tokamaks.

Typical exceptions, associated with the magnetic control, could be the loss of sensors resulting in a reduction in plasma

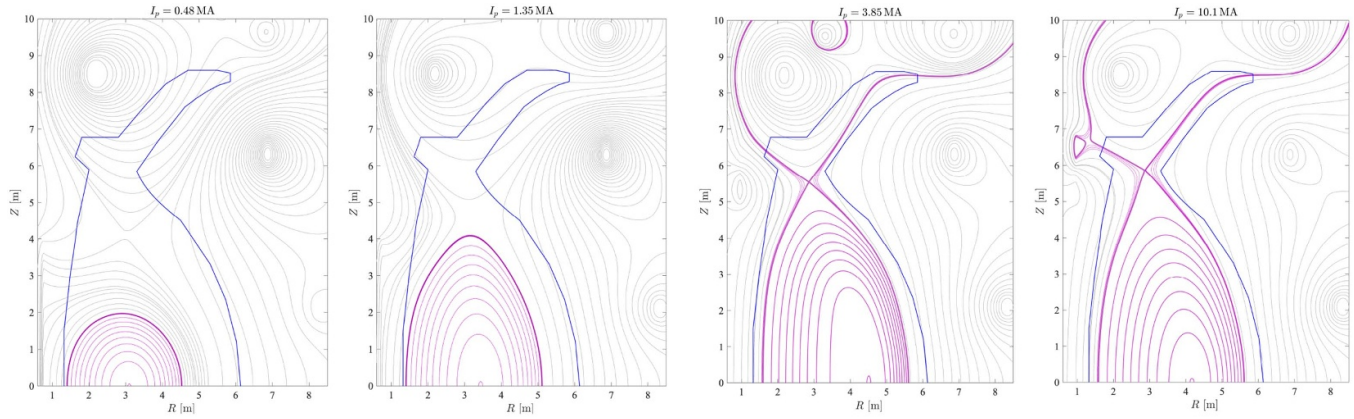

Figure 9. Equilibrium Evolution. As the equilibria are assumed to be up-down symmetric, only the upper half of the equilibria are shown.

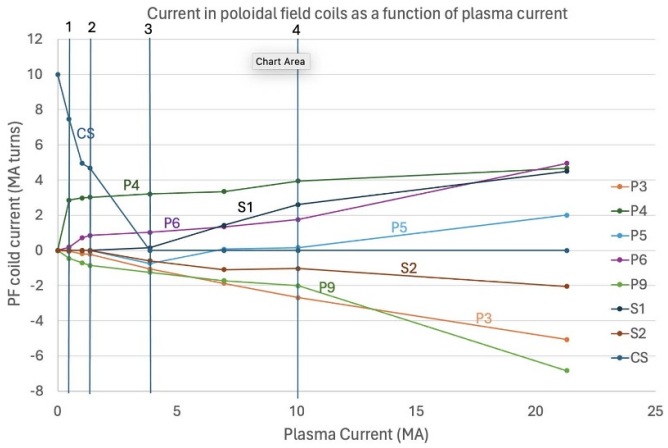

Figure 10. Poloidal Field Coil Current as a function of the plasma current during the ramp up phase. The coil current are shown in 'Ampere-Turns' as the number of turns per coil has not yet been fixed.

equilibrium reconstruction fidelity leading to the plasma shape no longer being as desired. This may result in over heating of plasma facing components. If real time model adaptation is able to restore a robust equilibrium reconstruction sufficiently quickly this would be a yellow exception. If the equilibrium reconstruction remains unreliable a termination trajectory has to be found which is tolerant to inaccuracies in equilibrium reconstruction.

4.3. Plasma vertical position control

As described above, the operation at high normalised pressure, required to achieve a high bootstrap current fraction, is facilitated by operating in a highly elongated plasma configuration. Such configurations are vertically unstable and require active stabilisation. The growth-rate associated with this instability increases with plasma elongation and it is reduced by the presence of toroidally conducting structures situated above and below the plasma. This reduction is due to currents induced in such structures when the plasma moves up or down, with the radial field generated by these induced currents counteracting the plasma motion. The closer the plasma is to the conducting structures and the higher the conductance of the structure, the more the vertical instability growth rate is reduced. During the vertical plasma movement, the currents induced in these structures flow in toroidally opposite directions above and below the plasma. Without such passive stabilising conducting structures, the vertical instability growth-rate would be so large that no active control system could maintain the vertical position. On the other hand, any conducting structures situated between the plasma and the active coils used by the vertical position control system will also slow down the magnetic field penetration from control coils to plasma, making active stabilisation more difficult. The passive stabilising structures must therefore be designed with care.

In figure 1, the magenta rectangles schematically represent a possible location for such stabilising structures. In this figure potential locations for the active stabilisation coils are shown in blue. Putting the active coils closer to the midplane than the passive stabilising structures limits the filtering effect on the active control fields due to the passive stabilising structures. The location of passive stabilising structures and active coils shown in figure 1 is just for illustration. If the passive conducting structures, shown in magenta in figure 1, were complete toroidal hoops, this is likely to prevent the initiation of the plasma using the central solenoid. In the case of complete hoops, the loop voltage generated by ramping the central solenoid current to initiate the plasma would induce significant current in these hoops as well as in the plasma. If the conductivity of these hoops is sufficiently large, the induced currents would create a large vertical field resulting in any nascent plasma potentially being elongated beyond the stabilising capability of the control system. To avoid this problem the conducting structures should not form full toroidal hoops. Instead, the upper and lower hoop should be connected in antiphase to prevent up-down toroidally symmetric currents from flowing while facilitating the generation of stabilising up-down anti-symmetric currents [37, 38]. In addition to shielding the plasma from magnetic fields generated by the active control coils, the passive conducting structures also shield any magnetic sensors from fields variations generated by plasma movements and as such this shielding also affects the sensing

of the plasma motion based on magnetic sensors. The appropriate design of conducting structures is therefore complex. The requirements for these structures derived from the need for effective vertical stabilisation can be summarised as follows:

The structures must:

(i) assure sufficiently low vertical instability growth rate,
(ii) not allow the induction of significant up-down symmetric currents during plasma formation,
(iii) permit the control field from the active stabilisation coils to penetrate to the plasma without excessive damping,
(iv) permit variations in the magnetic fields due to vertical motion of the plasma to reach magnetic sensors without excessive damping,

All conducting structures in the vicinity of the plasma or situated between control coils and the plasma will affect the vertical stabilisation. Most material near the plasma will be conducting as it is very hard to find any isolating material (other than vacuum) which can survive the neutron bombardment in this area. As such virtually all material situated within the vacuum vessel will be conducting with the design choice being between higher or lower conductivity. For the chosen reference plasma elongation (i) leads to the need to locate stabilising structures close to the plasma edge where first wall and blanket structures also need to be located. An integrated design selecting materials and conductive paths carefully is therefore required. The ERMES code has been developed to be able to compute the current induced in complex 3D structures due to variation in coil current or due to plasma instabilities, including vertical plasma movement [39]. Figure 11 shows an example of a calculation of such currents using ERMES. In this case the structures surrounding the plasma have been represented in a simplified manner by a number (8×32) of conducting modules. The conductivity for each module has been set to the same value, whereas the conductance in connections between modules have been selected either as highly conductive or with 0 conductivity. In this way conductive paths have been created. The 2nd and 7th rows have been connected to create an upper and a lower toroidal hoop structure, which is broken at one toroidal location. At this location the modules are connected poloidally to connect the two toroidal hoops in antiphase. This way of creating the antiphase connection is unlikely to be optimal, and the implementation just aims to be able to study the effect of the anti-symmetrically conducting hoops on the vertical stability and vertical position control. Other connections implemented between modules in rows 3–6 are related to RWM stabilisation and these will be discussed in the next section. The colours in figure 11 represent the intensity of currents induced in the connected modules due to a 1 MA m^{-2}, 0.1 Hz modulation of the current in the active vertical stabilisation coils. Note that the axisymmetric current in these coils induces currents both in the anti-series connected hoops dedicated to vertical stabilisation and in the smaller loops dedicated to RWM stabilisation. The figure illustrates the capability of ERMES to compute such non-toroidally symmetric induced currents. This code can also calculate the currents induced in the same structures due to plasma vertical motion and hence it can be used to determine how such structures affect the growth rate of the vertical instability. Based on this our toroidally symmetric vertical stabilisation models can be tuned to give similar behaviour without having to include ERMES in closed loop simulations.

Figure 12 shows a simulation of the closed loop controlled recovery of the vertical position after this position has been left to drift without active stabilisation to a position 10 cm above the unstable equilibrium position. This simulation has been done with a rigid body axisymmetric code (RZIp [40]) implemented in the 'Fiesta' [41, 42] equilibrium code. The toroidally symmetric conducting structures used in this simulation have been adapted to mimic the effect of the full 3D structures. The simulation has been carried out for 5 different maximum amplifier voltage limits (200 V, 400 V, 600 V, 800 V, 1 kV) as seen in the second box. In the simulation the pair of control coils used have been considered to have 2 turns each and to be situated inside the vacuum vessel, but outside the tritium breeding blanket as illustrated in figure 1. In addition to the damping and stabilising effects of the conducting structures, time constants of 2 ms and 5 ms have been considered to apply to measurements and amplifiers respectively. This figure, which shows an extreme example, indicates that, with the assumptions on plasma configuration and conducting structures, even significant vertical excursions would be recoverable, as long as the plasma is not further perturbed during the recovery phase. The design of the power supplies and coils required will be based on a trade-off between power supply voltage, maximum coil current and number of turns per coil, considering the expected range of plasma currents and instability growth rates.

To take advantage of the double null design choice, it is important the plasma configuration can be maintained close enough to perfect double null. In this connection it should be noted that a perfect double null configuration does not exist. One of the two nulls (x-points) will always be dominant. A measure for how close to perfect double null a plasma is, is the distance (dr_{sep}) between the separatrices associated with the two x-points as measured at the outboard midplane. The narrower the scrape-off layer outside the last closed flux surface is, the smaller dr_{sep} needs to be to take advantage of the double null and therefore the relevant measure of the achieved double null precision is dr_{sep}/λ, where λ is the scrape off layer decay length. Figure 13 shows a simplified illustration of the power distribution over the 4 divertor legs as a function of the deviation from perfect double null, as measured through the dr_{sep}/λ, parameter. This figure illustrates how the highest power striking any of the inner strike points could be reduced significantly by maintaining dr_{sep}/λ sufficiently small. More precise calculations of this power distribution using SOLPS are described in [12]. While showing somewhat lower power levels reaching the inner divertors, even when far from perfect double null, these calculations agree qualitatively

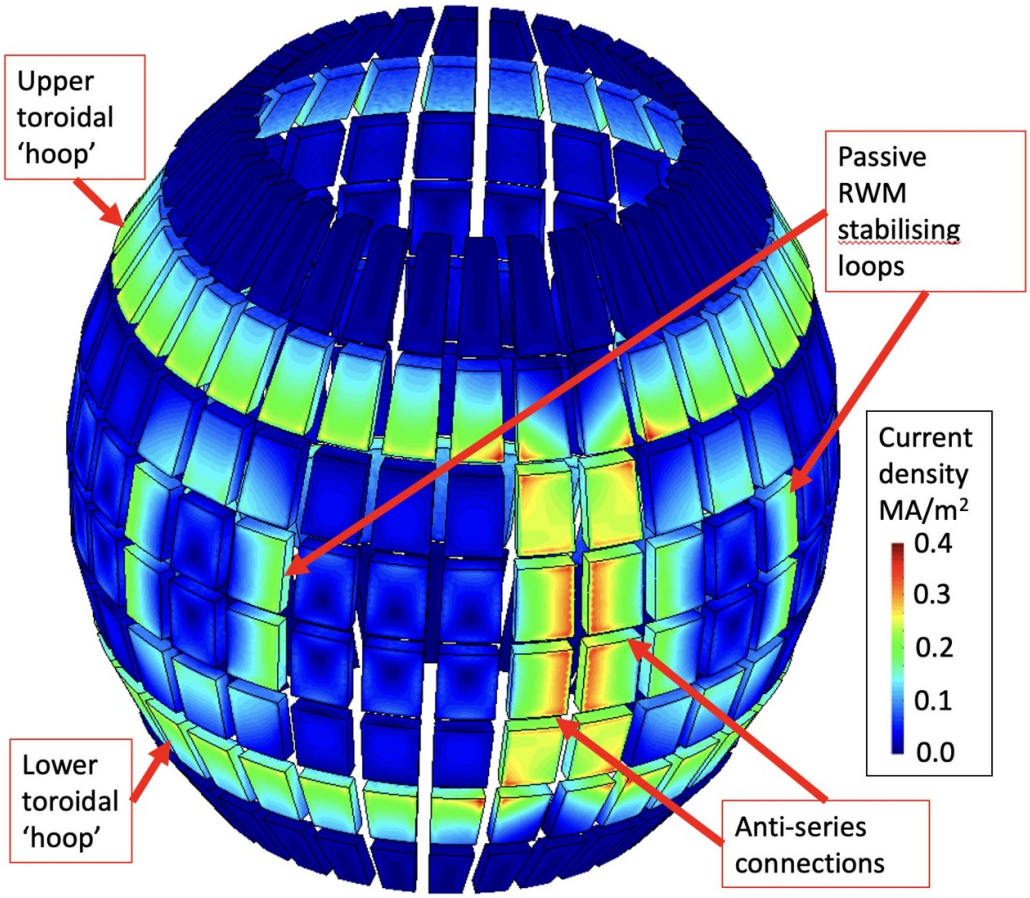

Figure 11. Induced currents in conducting structures surrounding the plasma due to a 1 MA m^{-2}, 0.1 Hz modulation coil current in the active vertical stabilisation coils. Conducting paths simulated by connecting neighbouring modules either with conducting or isolating material. (*i*) Up-down anti symmetric currents induced in rows 1 + 7 s. (*ii*) 2 smaller loops in rows 3–6 (with 2 other unseen toroidally opposite loops)—mainly aimed at RWM stabilisation.

with figure 13. Note that neither the calculations from [12], nor figure 13 includes the effect of particle drifts, which will break the up down symmetry [43]. If the up down symmetry is broken, the power could still be distributed equally between the upper and lower divertor by operating with a slight offset from pure double null. Doing this would, however, come at the cost of losing the additional reduction of the power towards the inner divertor legs which can be achieved when operating very close to pure double null. In addition the access to H-mode is likely to be significantly easier when close to a perfect double null [44]. For the STEP prototype λ is expected to be in the range of 1–2 mm in H-mode resulting in very stringent requirement on the control of dr_{sep}, which directly translates into stringent requirements on the control of the vertical position with dr_{sep} being approximately proportional to the vertical deviation Δz from the ideal position as follows:

$$dr_{sep} \cong 0.26\Delta z$$

This relation has been derived based on the preferred STEP plasma configuration using the FIESTA equilibrium code.

Due to the unstable nature of the vertical position, Δz will always fluctuate around 0 with the size of the fluctuation determined by the behaviour of the stabilising control loops. The control loop behaviour and through this the amplitude and frequency spectrum of these fluctuation will be determined by:

(i) the noise level in the measurement of the vertical velocity and position
(ii) the absolute accuracy of the measurement of the vertical position
(iii) the vertical instability growth-rate
(iv) the filtering due to conducting structures—affecting both control fields and position sensing
(v) the time constants and limits associated with power supplies and coils
(vi) the type and tuning of the controller.

Of these factors, the noise level and the time constants are likely to determine the ultimately achievable level of fluctuations during the steady state phase. While achieving very low fluctuation levels is required to take full advantage of the

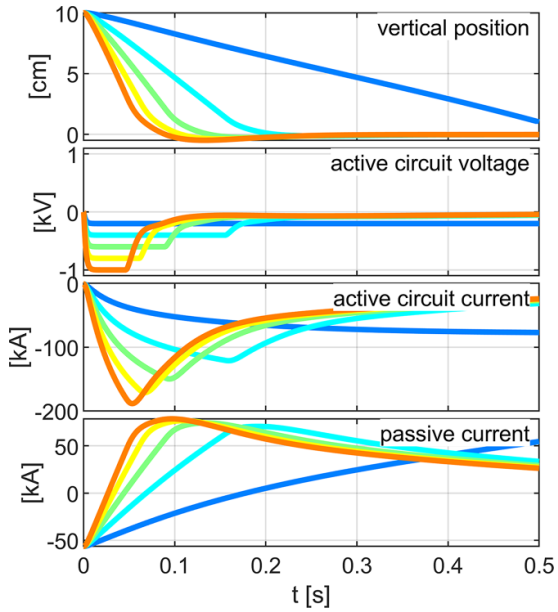

Figure 12. Closed loop simulation of the recovery from a 10 cm vertical excursion for different control voltage limits. (*i*) Vertical position evolution. (*ii*) Control voltages showing the limits. (*iii*) Control coils currents. (*iv*) Currents induced in passive coils.

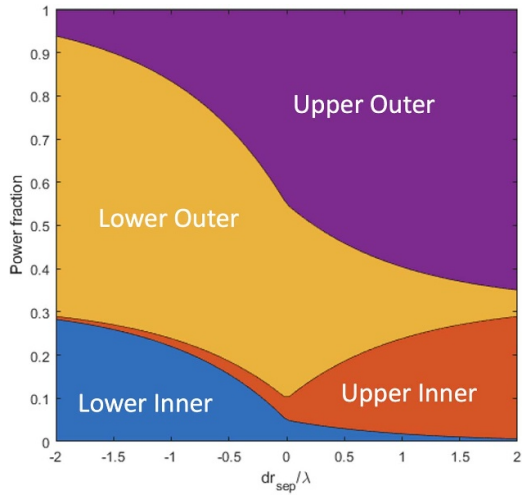

Figure 13. Idealised sketch illustrating the power distribution between the four strike-points when operating close to double-null, as a function of the deviation from pure double-null.

predicted additional reduction of heat loads towards the inner divertors, even without this additional gain, double null operation will still gain a factor of two in the average heat loads. In the case of significant fluctuations, it is desirable to assure that the frequency of such fluctuations is high with respect to the time constant associated with exit form H-mode and the thermal response of the relevant divertor components.

Though it is foreseen that the STEP prototype will operate without ELMs, it is worth noting that λ increases significantly during an ELM. This would mean that, even if the configuration is offset from a perfect double null, it is highly likely that any ELM energy will be directed mainly towards the outer strike-points, thereby reducing the risk of ELM damage to the more critical inner strike points. This consideration may influence the design of any systems implemented to eliminate ELM activity. Assuming RMP coils are used, this fact could be used to reduce the current in these coils until an ELM is detected and then increasing the current again to eliminate the ELM activity. In this way the threshold for ELM onset could be probed periodically, accepting the limited ELM activity entailed by such a scheme. A proposed arrangement of such RMP coils is shown in blue in figure 14.

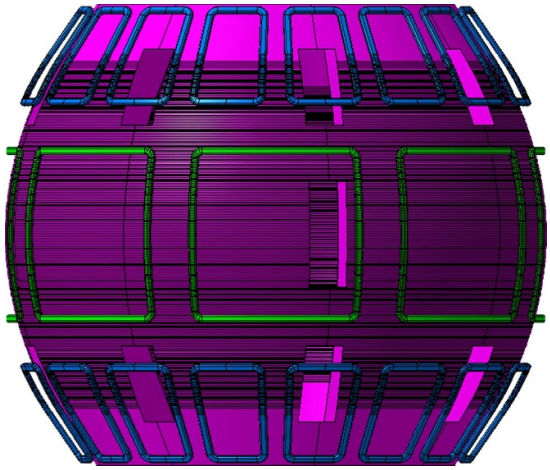

Figure 14. Possible arrangement of active in-vessel stabilisation coils for: (*i*) resistive wall mode control (green coils) (*ii*) ELM suppression (blue coils).

If the vertical stabilisation system fails, the plasma will accelerate up or down impacting the plasma facing components. Such an failure is described as a Vertical Displacement Event (VDE) and the ensuing disruption is one of the most severe events which the exception handling system have to prevent or mitigate. While loss of the measurment of the vertical velocity is likely to very rapidly lead to a VDE, the loss of measurement of the absolute position is not likely to cause a VDE. As a consequence the response of the exception handling system will be very different. The loss of vertical stabilisation will cleary bring the system into the red area in figure 5 and the response is likely to be the firing of the disruption mitigation system. If, on the other hand, the measurement of the absolute position becomes unreliable, while the velocity measurement is still available, the consequence would be an up-down heat load imbalance and the response is likely to be a reduction of the power reaching the divertors. This could be a combination of reducing the fusion power and increasing the plasma radiation. It is likely that this event would be considered to be in the orange area requiring a safe plasma termination.

4.4. Resistive wall mode control

The fusion gain is defined as $Q = P_{\mathrm{fus}}/P_{\mathrm{HCD}}$ where P_{fus} is the total fusion power and P_{HCD} is the external power injected into the plasma. High values of Q are required for net electricity production. For constant values of toroidal field

and plasma current, increasing the normalised beta β_N will be doubly helpful as an increase in pressure will increase the fusion power while an increased pressure gradient will increase the bootstrap current thereby reducing the required current drive power. Operation at high values of β_N may result in RWMs becoming unstable [20]. These modes have toroidal mode numbers $n \geqslant 1$, with $n = 1$ and $n = 2$ being most relevant here. Like the vertical instability, these modes are damped by currents induced in conducting structures surrounding the plasma (the resistive wall). Figure 11 shows a possible arrangement of conducting loops in rows 3–6 to form such a structure. An optimal design for these passive RWM structures is still under investigation. Note that the simulated currents seen in these loops in figure 11 are induced by variation of the current in the active vertical stabilisation coils rather than by variation of the current in the coils dedicated to RWM stabilisation. As for vertical stabilisation, it is important to select the conductivity distribution to achieve good mode stabilisation without excessive shielding of the fields from the active control coils. While an appropriate choice of conducting structures helps reducing the growth rate of the RWM, feedback stabilisation based on real time measurements of the mode, using active coils is likely to be required for sufficiently large β_N. A possible arrangement of such RWM coils is shown in green in figure 14. Further details of the work to design an active RWM stabilisation system, which will be very similar to the system needed for vertical stabilisation, can be found in [20].

A failure of the active RWM stabilisation system will, for sufficiently high β_N, nearly always lead to a disruption. If the power supply or other elements of this system approached a limit which could lead to a failure of the stabilisation, the exception handling system need to respond. The best response is likely to be to reduce β_N by reducing both density and fusion power. Such a reduction is likely to be reversible and hence in the yellow area in figure 5 as long as the RWM system can be returned to full capacity. If, on the other hand, the RWM system fails with little warning it is unlikely that β_N can be reduced in time. In this case a disruption is almost inevitable bringing the system into the red region of figure 5. In this case the exception handling system has to take action to mitigate the upcoming disruption.

4.5. Heat load management

In the double null design envisaged the outer divertor legs are brought to as large a major radius as possible in a 'Super-X' configuration in order to spread the heat load over as large an area as possible. This Super-X configuration has showed itself to be very effective on MAST-U [45]. Though the double null design aims at distributing the heat load between the divertor legs the power density at the strike points will still exceed the material limits without further mitigating measures. To distribute the heat load over a larger area the plan is to assure that most of the power leaves the plasma as electromagnetic radiation. Note that the power considered here does not include the 80% of the fusion power which leaves the plasma as 14 MeV neutrons and which will be absorbed in the tritium breeding blanket. If the radiation is induced solely in the divertor region, even with a good distribution of the heat load over the divertor surfaces, this load is still likely to exceed material limits. For this reason, a significant fraction of the power needs to be radiated from the plasma core. The core radiation will be increased by the injection of impurities. Ideally this radiation should originate from the plasma just inside the last closed flux surface while the radiation from more central parts of the plasma should be minimised. The radiating impurity should be selected such that it radiates at the lower plasma temperatures encountered near the edge, but not at the plasma temperatures encountered in the part of the plasma where the majority of the fusion power is produced. Xenon is currently being considered as a promising candidate for this. The absolute level of this core radiation has to be controlled to assure that sufficient power is conducted across the last closed flux surface to assure that the plasma remains in H-mode, while keeping the power conducted towards the divertors low enough.

It is desirable to minimise the required core radiation and hence the power that can be tolerated by the divertors should be maximised. To this end, almost all the power reaching the divertors should be radiated from the divertor plasma with only very cold plasma contacting the divertor surfaces. Achieving this, requires that the divertors are operated in the detached regime. Operating in this regime is also foreseen in other reactor scale tokamaks [42, 46, 47]. In this regime, volumetric radiation and interaction of the exhaust particles with neutrals and molecules, drives strong particle, energy, and momentum losses. This leads to a detachment of the ionisation region from divertor plates, drastically improving target conditions [48, 49]. When operating in this regime 'detachment' fronts are established along the divertor legs between the x-points and the strike points. The position of the detachment front should be controlled to strike a balance between core performance and divertor surface protection [8, 9].

The main actuator which can be used to control the detachment front position is matter (fuel and/or impurity) injection. The bandwidth of the detachment controller is determined by a combination of the delay and time constants associated with the matter injection and the vacuum pumping speed. Without sufficient pumping speed detachment control performance will be limited through a build-up of injected particles regardless of the delay/time constants associated with the matter injection.

The amount of impurity and/or fuelling required to maintain the detachment front at the desired position depends on the power flowing from the plasma via the scrape off layer towards the divertor. This power will vary due to plasma fluctuation and changes in external heating and fuel injection. A preliminary analysis has been carried out using JINTRAC to assess the influence of the injection of fuelling pellets on the power flowing towards the divertor. This is shown in figure 15. Both pure fuelling pellets and pellets doped with Xenon have been considered. It is seen that the power flowing to the divertor drops almost instantaneously. Evidently this drop would be too fast for the detachment controller to prevent some excursion in the position of the detachment front. For the fuelling pellets foreseen for the STEP prototype, the change in power is limited to

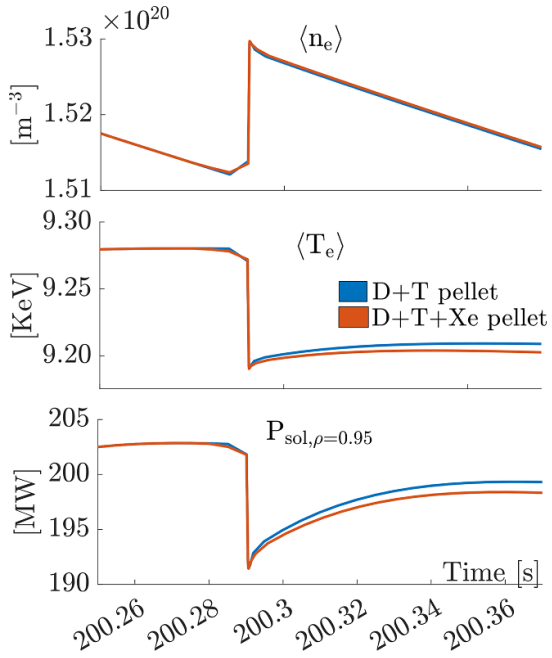

Figure 15. JINTRAC simulation of the effect of a single fuelling pellet injected into the core plasma. The curves shows the evolution of the (*i*) core density, (*ii*) core temperature (*iii*) Power flow into the scrape off layer which is conducted towards the divertors.

about 10%. This is significantly less than the expectations for ITER [50, 51] and DEMO, due to the smaller fuelling pellets required in a spherical tokamak. The detachment front movement due to such a change in power is expected to be acceptable and this means that the detachment controller would not need to respond to the injection of individual pellets and it would be ok for it to respond on a slower timescale to the evolution of the average power flowing towards the divertor. Note, however that this simulation does not consider the situation where pellets trigger brief confinement excursions which could lead to transient excursions in the power flowing towards the divertors.

Variations due to the fluctuation of the vertical position around the perfect double null will similarly result in variations in the power flow towards the divertors. This fluctuation could be much larger than the one associated with pellet fuelling as seen from figure 13 and hence the detachment controller may have to respond to such variations. Given the slow response expected from the detachment control actuators this may not be feasible if the vertical position variation is dominated by random noise induced fluctuations. If the fluctuations cannot be kept sufficiently small to allow the controller only to respond to the average power as described for pellet induced perturbations, it may be necessary to synchronise the vertical position and the divertor detachment matter injection oscillations, coupling the detachment controller to the vertical position controller. It would clearly be desirable to avoid this by guaranteeing sufficiently tight control of the vertical plasma position.

A more severe concern would be ELMs. It is clear that the detachment controller will not be able to respond to individual ELMs. To be able to tolerate a few sporadic ELMs, which will almost certainly burn through the detachment and strike the divertor surface, this surface has to be able to handle the associated energy pulses. If steady state ELMy operation is considered the size of such ELMs have to be small enough to avoid this 'burn through' and the ELM frequency has to be large enough for the detachment controller not to have to respond to individual ELMs.

Operation in double null means that there will be 4 strike points and the detachment state for each of these has to be controlled. It is not clear to which extent the upper and lower detachment controllers will be coupled, but it is almost certain that the controllers for the inner and outer strike points in the same divertor will be quite strongly coupled. A double null detachment control system is currently being prepared and tested on MAST-U [52].

As described in section 4.3 a failure to achieve good double null balancing is likely to lead to excessive power flowing towards one of the divertors, challenging the detachment controller and potentially leading to the loss of detachment. The response of the detachment controllers should be to increase the injection of neutrals and impurities in the divertor which receives the excessive power load. If a loss of detachment persist, some melting of the divertor surfaces is probably inevitable and in this case other controllers will be called upon to reduce the power flowing towards the divertors. This is a clear example of the exception handling involving several controllers. A persistent loss of detachment would reside in the orange area in figure 5 and the exception may rapidly escalate into the red region, with a disruption becoming unavoidable.

4.6. *Interaction of controllers*

The controller tasks described above are discussed separately, though they do not live in isolation. These controllers are all coupled with the coupling between the vertical and detachment controllers described above being an example. The overall plasma control system is in fact a large MIMO control system. As a fully decoupled control system would become very complicated, it is desirable to divide it into a number of separate MIMO and SISO (single input single output) controllers, between which coupling is sufficiently weak. If the coupling is weak enough, the coupling effects can be treated as disturbances. Though individual controllers can work without decoupling, it is likely that the controller performance can be improved by correct decoupling measures. The degree to which decoupling is implemented will therefore be a trade-off between simplification and performance [53–55].

5. Conclusion

The STEP project aims to build a nuclear fusion tokamak reactor which can operate at a sufficiently high value of the fusion gain Q to allow net electricity production. The ability to operate in steady state is, if not an absolute requirement, highly desirable for such a plant. Such operation requires that all the plasma current is generated through noninductive means. A

spherical tokamak has been selected as the most promising design since such a tokamak opens the way to operation with a high bootstrap current fraction, increasing Q by reducing the power required for external current drive. This choice leads to more demanding plasma control challenges and therefore the design of the plasma control system has started early in the STEP conceptual design phase. This allows the requirements of the control system to be well defined early in the project. These control system requirements can then be translated into requirements both for the actuators used by the control system and for the general machine design. The plasma control design efforts have initially focused on the most critical control challenges: (i) vertical stabilisation, (ii) RWM stabilisation, (iii) control of the plasma profiles, (iv) detachment control.

It is essential to guarantee that it is possible to stabilise the highly elongated, high β_N plasmas needed for high bootstrap current operation. Achieving robust stability requires good control of the plasma current radial distribution and careful design of the conductivity distribution in the structures surrounding the plasma. This, in turn, translates into requirements for: (i) amplifiers and coils used for plasma vertical position and RMP stabilisation and general plasma shape and position control, (ii) heating, current drive and fuelling systems which allow control of the strongly coupled temperature, density and current profiles. Similarly, the control of the detachment feeds into requirements for matter injection systems and vacuum pumping.

The individual controllers have to operate in conjunction to assure that the plasma is brought from plasma initiation to steady state full power operation, before being terminated safely. The controllers have to assure that the plasma follows robust trajectories within the relevant multidimensional parameter space. During the steady state phase, the controllers have to assure that the plasma remains in the desired state and, more specifically, that the fusion power produced is controlled to the desired value.

The hostile environment expected in the STEP prototype is likely to result in a significant limitation of the measurements available for the controllers. To deal with this and the expected slow plasma evolution it is expected that plasma control will rely heavily on model-based control. Comparing the modelled parameter evolutions with measurements can also help identifying exceptions which need special handling by the plasma control system. Handling such exceptions is one of the most important tasks for the plasma control system, with the main goal being to avoid exceptions escalating to the level where a disruption is unavoidable.

The STEP plasma control design work relies heavily on having appropriate models around which controllers can be designed. A number of separate models already exist, and they have been used in the work presented here. These models are, however, not well integrated. The development of good, integrated, 'control oriented' models is therefore a high priority task. Such models should capture the main system dynamics and non-linearities, but they do not have to be derived directly from the underlying physics. Surrogate models derived from data measured on existing devices or data generated by physics based on higher fidelity codes are likely to play an important role. A hierarchy of codes of different fidelity levels will be required, with faster than real time capable codes being required for model predictive control, while more accurate, though not too slow, codes will be required for the development of the control algorithms. These codes should themselves be validated against high fidelity codes. Having reasonably rapid validated integrated models for the entire plasma control system will allow the requirements for actuators, sensors etc. to be derived based on the need for robust control.

ORCID iDs

M. Lennholm https://orcid.org/0000-0002-3444-3999
S. Aleiferis https://orcid.org/0000-0001-7529-470X
O.P. Bardsley https://orcid.org/0000-0002-1525-675X
M. van Berkel https://orcid.org/0000-0001-6574-3823
F.J. Casson https://orcid.org/0000-0001-5371-5876
S.S. Henderson https://orcid.org/0000-0002-8886-1256
A. Hudoba https://orcid.org/0000-0002-8197-7977
B. Kool https://orcid.org/0000-0002-1636-7400
H. Meyer https://orcid.org/0000-0002-5565-1199
R. Otin https://orcid.org/0000-0002-3053-2695
G. Xia https://orcid.org/0009-0009-4236-077X

References

[1] Wilson H. *et al* 2020 STEP—on the pathway to fusion commercialization *Commercialising Fusion Energy* (IOP Publishing) (https://doi.org/10.1088/978-0-7503-2719-0ch8)
[2] Meyer H. *et al* 2023 The plasma scenarios for the spherical tokamak for energy production (STEP) and their technical implications *Proc. 29th IAEA Fusion Energy Conf.* (*London, UK, 16–21 October 2023*) (available at: https://conferences.iaea.org/event/316/papers/28532/files/10411-HMeyer_IAEA2023_paper_v2.pdf)
[3] Muldrew S.I. *et al* 2023 *Fusion Eng. Des.* **201** 114238
[4] Gerhardt S.P. *et al* 2011 *Nucl. Fusion* **51** 073031
[5] Kessel C.E. *et al* 1994 *Nucl. Fusion* **34** 1221
[6] Lennholm M. *et al* 2023 Plasma control for the STEP power plant *Proc. 29th IAEA Fusion Energy Conf.* (*London, UK, 16–21 October 2023*) (available at: https://conferences.iaea.org/event/316/papers/28230/files/10650-Lennholm_IAEA_FEC_2023_Paper_230919a.pdf)
[7] Meschini S., Ferry S.E., Delaporte-Mathurin R. and Whyte D.G. 2023 *Nucl. Fusion* **63** 126005
[8] Henderson S. *et al* 2023 *Nucl. Fusion* **63** 086024
[9] Henderson S. *et al* 2023 Divertor detachment and reattachment with mixed impurity seeding on ASDEX Upgrade and JET *Proc. 29th IAEA Fusion Energy Conf.* (*London, UK, 16–21 October 2023*) (available at: https://conferences.iaea.org/event/316/papers/28036/files/10632-IAEA_FEC_2023.pdf)
[10] Hudoba A., Newton S., Voss G., Cunningham G. and Henderson S. 2023 *Nucl. Mater. Energy* **35** 101410
[11] Hudoba A., Cunningham G. and Bakes S. 2023 *Fusion Eng. Des.* **19** 113704
[12] Osawa R.T., Moulton D., Newton S.L., Henderson S.S., Lipschultz B. and Hudoba A. 2023 *Nucl. Fusion* **63** 076032
[13] Brunner D., Kuang A.Q., LaBombard B. and Terry J.L. 2018 *Nucl. Fusion* **58** 076010
[14] Meyer H. *et al* 2005 *Plasma Phys. Control. Fusion* **47** 843
[15] Joffrin E. *et al* 2019 *Nucl. Fusion* **59** 112021
[16] Garzotti L. *et al* 2019 *Nucl. Fusion* **59** 076037

[17] Bernert M. *et al* 2021 *Nucl. Fusion* **61** 024001
[18] Ryan D.A., Dunne M., Kirk A., Saarelma S., Suttrop W., Ham C., Liu Y.Q. and Willensdorfer M. 2019 *Plasma Phys. Control. Fusion* **61** 095010
[19] Loarte A. *et al* 2014 *Nucl. Fusion* **54** 033007
[20] Xia G., Liu Y., Hender T.C., McClements K.G., Trier E. and Tholerus E. 2023 *Nucl. Fusion* **63** 026021
[21] Henderson M. *et al* 2023 The concept design of the STEP HCD system *Proc. 29th IAEA Fusion Energy Conf.* (*London, UK*, *16–21 October 2023*) (available at: https://conferences.iaea.org/event/316/papers/28688/files/10954-IAEA-Mark-Manuscript-clean.pdf)
[22] Freethy S. *et al* 2023 The STEP microwave heating and current drive system *Proc. 29th IAEA Fusion Energy Conf.* (*London, UK*, *16–21 October 2023*) (available at: https://conferences.iaea.org/event/316/papers/28139/files/11041-Manuscript.pdf)
[23] Snipes J.A. *et al* 2021 *Nucl. Fusion* **61** 106036
[24] Biel W. *et al* 2022 *Fusion Eng. Des.* **179** 113122
[25] Mitchell J., Parrott A., Casson F.J., Eriksson F.E., Koechl F., Lennholm M., Bardsley O., Meyer H.F. and Aleiferis S. 2023 *Fusion Eng. Des.* **92** 113777
[26] Blanken T.C., Felici F., Galperti C., Vu N.M.T., Kong M., Sauter O. and de Baar M.R. 2019 *Nucl. Fusion* **59** 026017
[27] Fil A. *et al* 2023 Disruption runaway electron generation and mitigation in the spherical tokamak for energy production *Proc. 29th IAEA Fusion Energy Conf.* (*London, UK*, *16–21 October 2023*) (available at: https://conferences.iaea.org/event/316/papers/28547/files/10745-AFIL%20-%20IAEA%20FEC%202023%20Paper%20-%20Final.pdf)
[28] Tholerus E. *et al* 2024 Flat-top plasma operational space of the STEP power plant *Nucl. Fusion* submitted
[29] Anand H., Bardsley O., Humphreys D., Lennholm M., Welander A., Xing Z., Barr J., Walker M., Mitchell J. and Meyer H. 2023 *Fusion Eng. Des.* **194** 113724
[30] Felici F., Citrin J., Teplukhina A.A., Redondo J., Bourdelle C., Imbeaux F. and Sauter O. 2018 *Nucl. Fusion* **58** 096006
[31] Mitchell J. *et al* 2024 in preparation
[32] Romanelli M. *et al* 2014 *Plasma Fusion Res.* **9** 3403023
[33] Eriksson F. *et al* 2022 Current ramp-up modelling for STEP *48th EPS Conf. on Plasma Physics* (*Maastricht, the Netherlands*, *27 June–1 July 2022*) (available at: https://indico.fusenet.eu/event/28/contributions/260/attachments/274/1134/FEriksson_Paper_EPS22.pdf)
[34] Greenwald M. 2002 *Plasma Phys. Control. Fusion* **44** R27
[35] Moret J.-M., Duval B.P., Le H.B., Coda S., Felici F. and Reimerdes H. 2015 *Fusion Eng. Des.* **91** 1–15
[36] Ambrosino R., Acampora E., Albanese R., Castaldo A., Di Marzo V., Maviglia F., Portone A. and Zohm H. 2023 *Fusion Eng. Des.* **197** 114029
[37] Vorpahl C., Suttrop W., Ebner M., Streibl B. and Zohm H. 2013 *Fusion Eng. Des.* **88** 537–40
[38] Gruber O., Kaufmann M., Köppendörfer W., Lackner K. and Neuhauser J. 1984 *J. Nucl. Mater.* **121** 407–14
[39] Otin R. 2013 *Comput. Phys. Commun.* **184** 2588–95
[40] Coutlis A., Bandyopadhyay I., Lister J.B., Vyas P., Albanese R., Limebeer D.J.N., Villone F. and Wainwright J.P. 1999 *Nucl. Fusion* **39** 663
[41] Cunningham G. 2013 *Fusion Eng. Des.* **88** 3238–47
[42] Wenninger R.P. *et al* 2014 DEMO divertor limitations during and in between ELMs *Nucl. Fusion* **54** 114003
[43] Petrie T.W. *et al* 2008 *Nucl. Fusion* **48** 045010
[44] Meyer H. *et al* 2006 *Nucl. Fusion* **46** 64
[45] Wijkamp T.A. *et al* 2023 *Nucl. Fusion* **63** 056003
[46] Pitts R.A. *et al* 2019 Physics basis for the first ITER tungsten divertor *Nucl. Mater. Energy* **20** 100696
[47] Subba F., Coster D.P., Moscheni M. and Siccinio M. 2021 SOLPS-ITER modeling of divertor scenarios for EU-DEMO *Nucl. Fusion* **61** 106013
[48] Verhaegh K. *et al* 2023 *Nucl. Fusion* **63** 126023
[49] Verhaegh K. *et al* 2023 *Nucl. Fusion* **63** 016014
[50] Wiesen S., Köchl F., Belo P., Kotov V., Loarte A., Parail V., Corrigan G., Garzotti L. and Harting D. 2017 *Nucl. Fusion* **57** 076020
[51] Garzotti L. *et al* 2019 *Nucl. Fusion* **59** 026006
[52] Kool B. *et al* 2024 in preparation
[53] Piron L. *et al* 2021 *Fusion Eng. Des.* **166** 112305
[54] Lennholm M. *et al* 2017 *Fusion Eng. Des.* **23** 535–40
[55] Koenders J.T.W., Perek A., Galperti C., Duval B.P., Février O., Theiler C. and van Berkel M. 2023 *Nucl. Fusion* **63** 106007