PAPER • OPEN ACCESS

Decoupled magnetic control of spherical tokamak divertors via vacuum harmonic constraints

Decoupled magnetic control of spherical tokamak divertors via vacuum harmonic constraints

O P Bardsley[1,*], **J L Baker[2]** and **C Vincent[1]**

[1] United Kingdom Atomic Energy Authority, Culham Campus, Abingdon OX14 3DB, United Kingdom
[2] York Plasma Institute, Department of Physics, University of York, Heslington YO10 5DD, United Kingdom

E-mail: oliver.bardsley@ukaea.uk, john.lloyd.baker@york.ac.uk and charles.vincent@ukaea.uk

Abstract
Power exhaust is a critical challenge for spherical tokamak reactors, making the design, optimisation and control of advanced divertor configurations crucial. These tasks are greatly simplified if the poloidal magnetic fields in the core and divertor regions can be varied independently. We present a novel method which facilitates decoupling of the core plasma equilibrium from the divertor geometry optimisation and control, using vacuum spherical harmonic (SH) constraints. This has the advantage that it avoids iterative solution of the Grad–Shafranov equation, making it easy to use, rapid and reliable. By comparing a large number of MAST-U equilibrium reconstructions against their approximations using SHs, a small number ($\sim$4) of harmonics is found to be sufficient to closely reproduce the plasma boundary shape. We show experimentally that poloidal field changes designed to leave harmonics unaffected indeed have no effect on the core plasma shape. When augmented with divertor geometry constraints, this approach gives a powerful tool for creating advanced magnetic configurations, and its simplicity brings improvements in speed and robustness when solving coil position optimisation problems. We discuss the clear benefits to real-time feedback control, feed-forward scenario design and coilset optimisation with a view to future reactors.

Keywords: spherical tokamaks, MHD equilibrium, plasma control, divertor, spherical harmonics, coilset optimisation

1. Introduction

The spherical tokamak (ST) is attracting significant attention as a source of fusion energy [1–6] due to its potential for high plasma energy density and therefore reduced size and cost [7], made possible by recent advances in magnet technology [8]. However, the compactness of STs exacerbates the challenge of exhaust power handling, necessitating advanced divertor magnetic configurations to keep plasma facing components within material limits [2, 9]. This necessitates

- Devising robust real-time control algorithms for divertor magnetic control in present and future devices;
- Testing candidate divertor magnetic geometries in present-day experimental devices (e.g. MAST-U);
- Designing poloidal field (PF) coil positions and currents capable of producing the scenarios required for a future power plant.

A commonality here is the desire to keep a fixed core plasma shape (to decouple shape and divertor controllers, to ensure

---

* Author to whom any correspondence should be addressed.

fair comparison, or to produce a particular burning plasma solution, respectively), whilst the divertor geometry is independently specified or controlled—a concept we henceforth refer to as just 'decoupling'. Because each PF coil affects both the core plasma and divertor magnetic geometries, controlling just the latter requires simultaneous adjustment of all the PF coil currents.

Existing inverse free-boundary equilibrium codes [10–15] generally solve this problem by minimising the distance of the plasma boundary from a set of control points, which may extend into the divertor region. The second-order non-linear Grad–Shafranov equation (GSE), which determines the plasma shape, is iteratively solved concurrently with the optimisation. This approach—which is compulsory for creating the core plasma shape—is well-understood, but when combining with divertor geometry optimisation can be challenging to converge, requiring user expertise and time to appropriately formulate; this drawback is particularly acute when allowing coil positions to vary as well.

Here we demonstrate a novel decoupling approach which does not require iteration of the GSE. We make use of the fact that the core plasma shape is maintained as long as the vacuum magnetic field (i.e. that due to the PF coils) over the region of space it occupies remains unchanged, for a given internal current profile. Furthermore, for a ST the vacuum field is well-approximated by a spherical harmonic (SH) expansion truncated at just a few terms. Constraining these terms whilst the PF currents are varied provides a simple way to achieve decoupling without appeal to the GSE.

As complete, orthogonal solutions to the Laplace equation, SHs are widely employed in physics applications. For example, to model the effect of ferromagnetic walls on tokamak stability [16], in inertial confinement fusion for three-dimensional reconstructions [17] and compression models of radiation [18], for magnetic resonance imaging [19], and representing Earth's magnetic field [20, 21].

The objective of this paper is to describe and validate the method, and to demonstrate its power in addressing the three key challenges posed above. In section 2, we outline how SHs may be used to formulate simple constraints which replace solution or linearisation of the GSE, yet still achieve decoupling. Our results consist of numerical verification in section 3.1, experimental verification through control application on MAST-U in section 3.2, examples of scenario design in section 3.3, and coil position optimisation in section 3.4. In section 4 we discuss how this approach may be exploited to meet the divertor design and control needs of present and future STs.

2. Methods

We begin by recalling key aspects of tokamak equilibrium, then outline how vacuum SH constraints can be used to replace iteration of the GSE when decoupling divertor geometry from the core plasma.

In cylindrical polar co-ordinates $[R, \phi, Z]$, the poloidal magnetic flux density $\boldsymbol{B}_p(R, Z)$ for an axisymmetric tokamak equilibrium is found by integrating Ampère's law (i.e. the GSE),

$$\boldsymbol{e}_\phi \cdot \nabla^2 (A\boldsymbol{e}_\phi) = -\mu_0 j, \tag{1}$$

where $A(R, Z)\boldsymbol{e}_\phi$, the vector potential for $\boldsymbol{B}_p$, is related to the poloidal flux $\psi = AR$, $\boldsymbol{e}_\phi$ is the toroidal unit vector, $j(R, Z)$ is the toroidal current density, and μ_0 is the vacuum permeability.

The current density, and therefore A, has contributions due the plasma and PF coils, $j = j_{\text{pla}} + j_{\text{PF}}$. (For clarity we neglect passive structure currents here, but their contribution is similar to the PF coils.) j_{pla} is non-zero only in the region of space within the core plasma boundary Ω_{pla}, taking the form

$$j_{\text{pla}} = Rp' + \frac{ff'}{\mu_0 R}, \tag{2}$$

$p(\psi)$ being the plasma pressure and $f(\psi)$ the toroidal field function. j_{PF} can be represented by a set of δ-functions representing PF coil filaments.

If p' and ff' are known, the free-boundary equilibrium problem is to solve this system with the PF coil currents as inputs and the plasma boundary as an output, or vice versa, or some hybrid of the two. It is non-linear not only because (2) depends on ψ, but also because Ω_{pla} is unknown; this difficulty means solving the GSE necessitates an iterative method with the potential to be time-consuming and numerically unstable.

Now suppose that this iteration has already been completed once, meaning Ω_{pla} and the PF currents have been determined for some base case; this could be an experimental equilibrium reconstruction, or an inverse problem for the PF currents to give a desired core plasma shape. (Note this solution does not have to obey the constraints of—or even use the same coilset as—the following calculation.) We wish to alter the PF currents in some way that preserves this core plasma shape Ω_{pla} and current distribution j_{pla}, but decouples the divertor magnetic geometry. It is key to realise that this wish will be satisfied if the field due to the PF coils (i.e. A_{PF}) within the domain Ω_{pla} does not change. The only way to keep $A_{\text{PF}}|_{\Omega_{\text{pla}}}$ strictly invariant is to use exactly the same PF currents. However, the geometry of an ST lends itself naturally to representing $A_{\text{PF}}|_{\Omega_{\text{pla}}}$ as a series of SHs about the origin, and we assert that keeping just the first few terms of this series constant will in practice prove sufficient. So long as these coefficients are unchanged when choosing new PF currents, the core plasma shape will not change appreciably.

SH coefficients are easily related to the coilset geometry and PF currents. Introducing spherical polar co-ordinates $[r, \theta, \phi]$ so $R = r\sin\theta$ and $Z = r\cos\theta$, we have [22]

$$A_{\text{PF}} = \sum_{\ell=1}^{\infty} A_{\text{PF}}^\ell \left(\frac{r}{r_0}\right)^\ell \frac{P_\ell^1(\cos\theta)}{\sqrt{\ell(\ell+1)}}, \tag{3}$$

where $P_\ell^1(\cdot)$ is the associated Legendre polynomial of degree ℓ and order 1, and r_0 is a scaling length. For a set of toroidal

filaments '*f*' (making up a coil or whole coilset, say), the SH coefficients are given by

$$A_{\mathrm{PF}}^{\ell} = \sum_f \frac{\mu_0}{2} I_f \sin\theta_f \left(\frac{r_0}{r_f}\right)^{\ell} \frac{P_{\ell}^1(\cos\theta_f)}{\sqrt{\ell(\ell+1)}}, \qquad (4)$$

with I_f being the current in each filament. We calculate the A_{PF}^{ℓ} coefficients for the base case using (4), then require that the first $\ell_{\max}$ of them are fixed as the PF coil currents (and possibly positions) vary, again using (4). This gives $\ell_{\max}$ equality constraints (hereafter 'SH constraints') on the problem which entirely replace the GSE, and are linear in the free variables (I_f) for fixed coil positions.

It is important to note that (3) and (4) are only valid for locations within an imaginary sphere about the origin which does not contain any filaments, i.e. $r < \min(r_f)$. To use the method in its primitive form, this sphere must therefore contain the whole plasma volume, and hence only PF circuits with all $r_f > \max\left(r_{\partial\Omega_{\mathrm{pla}}}\right)$ may be free. This geometrical consideration limits the practicality to STs, and furthermore means that centre-column PF currents must be held fixed, or otherwise compensated for in the optimisation (see appendix).

SH constraints can be fully exploited by combining them with divertor geometry constraints. We limit discussion here to linear constraints, such as those on poloidal flux and field; for example, suppose we wish a point *P* to lie on the separatrix ('sep'), a desire we can write as

$$\psi|_P = \psi_{\mathrm{PF}}|_P + \psi_{\mathrm{pla}}|_P = \psi_{\mathrm{sep}}. \qquad (5)$$

This gives a simple constraint on $\psi_{\mathrm{PF}}|_P$, which is merely a weighted sum of the PF currents. This is only possible because the SH constraints ensure a fixed core plasma, meaning both $\psi_{\mathrm{pla}}|_P$ and ψ_{sep} do not change. Note that the plasma shape (core and divertor) is fixed here via exact constraints, rather than as a contribution to an objective function to be minimised. These linear constraints can be written in matrix form

$$\mathsf{G}\boldsymbol{x} = \boldsymbol{c}, \qquad (6)$$

where $\boldsymbol{x}$ is a column vector containing the circuit currents, and the matrix G and column vector $\boldsymbol{c}$ are used to encode SH constraints (i.e. (4) for the first few values of ℓ, with A_{PF}^{ℓ} known from the base case), divertor constraints (e.g. (5), for which G would contain a row of Green's function coefficients and $\boldsymbol{c}$ would contain $\psi_{\mathrm{sep}} - \psi_{\mathrm{pla}}|_P$), and any other linear constraints (such a fixing the current in a particular circuit). As long as (6) is satisfied, we are free to change $\boldsymbol{x}$ however we wish; assuming fewer constraints than degrees of freedom, its general solution is

$$\boldsymbol{x} = \mathsf{G}^+\boldsymbol{c} + \mathsf{N}\boldsymbol{w}, \qquad (7)$$

where '$+$' denotes the pseudoinverse, N is a matrix whose columns span the nullspace of G, and $\boldsymbol{w}$ is a vector of weights encapsulating the remaining degrees of freedom, which can be used to vary the PF currents.

3. Results

3.1. Numerical verification

Our first task is to verify numerically that constraining a small number of vacuum SHs successfully preserves the core plasma shape. We take the EFIT++ reconstruction [23–25] of a MAST-U plasma and PF currents, recalculate the GSE equilibrium using the free-boundary code Fiesta [26], then use (4) to calculate the SH coefficients due to the (non-centre-column) PF circuits. Keeping centre-column currents fixed, we replace the field due to the other PF circuits with the series (3) truncated at degree $\ell_{\max}$, then recalculate the GSE equilibrium once again. We assess the similarity between the two equilibria by evaluating a fit metric

$$\eta\left(\Omega_{\mathrm{pla}}^1, \Omega_{\mathrm{pla}}^2\right) := \frac{\mathrm{XOR}\left(\Omega_{\mathrm{pla}}^1, \Omega_{\mathrm{pla}}^2\right)}{\Omega_{\mathrm{pla}}^1 + \Omega_{\mathrm{pla}}^2}, \qquad (8)$$

i.e. the area of the poloidal plane contained by only one plasma cross-section, divided by the sum of the areas—so a smaller η implies a better fit. This processes is repeated for 640 different plasmas (picked at random from MAST-U shots 43761–46799, times 0.1–1 s, and with plasma current over 100 kA), to give a statistical distribution for η as a function of $\ell_{\max}$.

Figure 1(a) shows how disparity between the base equilibrium and its approximation using SHs robustly decreases as $\ell_{\max}$ increases; an example case, which lies close to the median, is shown in figure 1(b) to give intuition for interpretation of η. We see that a visual match to the plasma shape is obtained for $\eta \lesssim 10^{-2}$, and this is satisfied by the majority of cases when using 5 odd SHs ($\ell_{\max} = 9$). (Even SHs have up-down antisymmetric ψ, so are practically zero for the symmetric double-null plasmas of MAST-U.) In practice we find that 4 odd SHs are sufficient when combined with divertor constraints, and use this in the examples which follow, checking that $\eta < 10^{-2}$ throughout.

3.2. Experimental verification and control application

We now verify the use of SH constraints experimentally, and in so doing illustrate their utility to plasma control and suitability for real-time application.

Our approach is to design virtual circuits (VCs) for decoupling control of quantities of interest from the core plasma. A VC $\delta\boldsymbol{x}$ is simply a set of weights applied to PF currents; the MAST-U plasma control system architecture is designed to facilitate their flexible implementation [28]. Suppose the change in a control quantity q is linearly related to the PF currents by $\delta q = \boldsymbol{g}^t \delta\boldsymbol{x}$ for some vector $\boldsymbol{g}$ (we give a couple of examples below). So that SH and other constraints (6) remain satisfied, the VC must satisfy $\mathsf{G}\delta\boldsymbol{x} = \boldsymbol{0}$, and therefore (7) becomes $\delta\boldsymbol{x} = \mathsf{N}\delta\boldsymbol{w}$. The optimal choice of VC is that for which $\delta\boldsymbol{w}$ has no component in any direction which *does not* change q, i.e.

$$\delta\boldsymbol{x} \propto \mathsf{N}\mathsf{N}^t\boldsymbol{g}. \qquad (9)$$
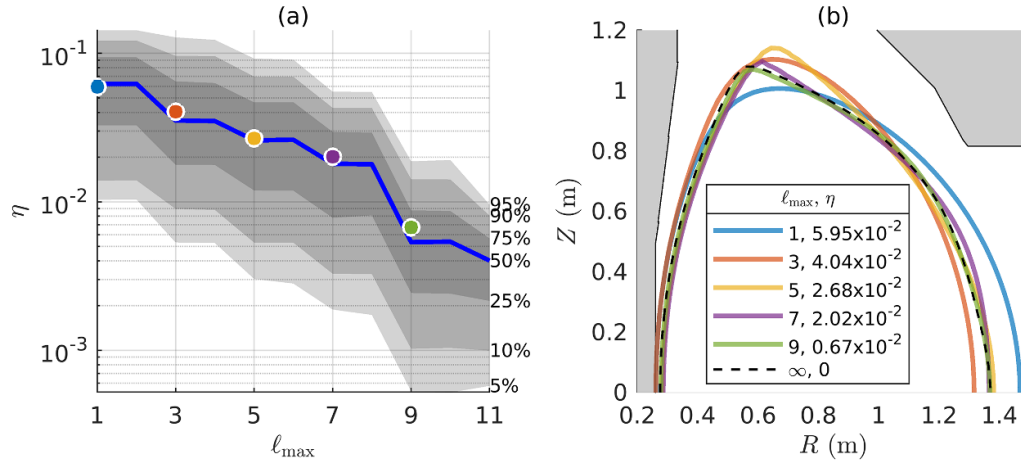
Figure 1. (a) Percentiles of the distribution of the fit metric η across 640 MAST-U equilibria as a function of the number of SHs used to approximate the vacuum field; (b) example boundary shapes and corresponding fit metrics for a single case (shot 44626 at 390 ms), corresponding to markers in (a). Data can be found at [27].

We now present the application this approach to designing VCs in three experiments: targeted reduction of a PF current in feed-forward, and two super-X sweeps—one in feed-forward and one in feedback.

3.2.1. Circuit current reduction.

The first problem we consider is that of reducing a PF current which is close to its limit and redistributing it amongst the other circuits, without compromising the plasma shape. We choose to target the 'DP' coils, which are in the divertor noses and tend to operate at a large negative current; the ***g*** in this case is therefore zeros, except a 1 at the index corresponding to DP. Our constraints on the VC are that (i) the first 4 odd SHs it produces are all zero, (ii) its contribution to poloidal flux at a point in the divertor throat is zero, (iii) the currents in centre-column (P1, PC, PX) and vertical control (P6) circuits are zero. Forming G and applying (9) produces the VC shown in figure 2(d).

To experimentally validate this we take MAST-U shot 48709—a 750 kA Ohmic plasma with conventional divertor—as a reference, and add a linear ramp in the SH-derived VC from 300–400 ms to effect a +500 A change (13% reduction in magnitude) in the DP current; figure 2(a) shows the measured DP current for the new shot (48713) has indeed matched this request. Despite this substantive change, there is no observable effect on the core plasma shape. In figure 2(b) we show the fit metric (8) comparing the two shots, which remains well below 10^{-2} throughout the VC ramp. In itself this is not remarkable, because feedback control for the core plasma shape was turned on for both shots. However, the drive terms for these shape controllers, plotted in figure 2(c), were very similar, meaning they are not having to do any work to correct for the effect of the new VC.

To give a sense for how the VC works we plot its flux surfaces in 2(e); the poloidal flux due to the VC is very small over the core plasma chamber and at the 'throat' control point (red cross). To reiterate the similarity between the equilibria and show that the divertor leg has also been unaffected by this VC,

we also compare the EFIT-reconstructed flux surfaces at the end of the ramp.

3.2.2. Super-X sweep.

Our second VC is designed to control the strike point location across the whole divertor, between a conventional ($\sim$0.8 m radius) and super-X ($\sim$1.4 m). The poloidal flux at a point *P* roughly in the middle (see red target in figure 3(c)) is therefore used for *q*, so ***g*** contains the flux at *P* due to each circuit. We simply constrain to zero (i) 4 odd SHs, (ii) current in P1/PC/PX/P6, from which (9) gives the VC in figure 3(a) with flux surfaces in figure 3(c).

First we apply this VC in feed-forward. We add a linear ramp in its amplitude over 350–600 ms onto a reference shot (48775), which produces a full strike point sweep from conventional to super-X (shot 49067) with a single control action—see figure 3. The match to reference at 600 ms is $\eta = 5.8 \times 10^{-3}$.

Next we achieve the same results in feedback, making use of the real-time strike point estimation capabilities of MAST-U [29–31] to control its radial location, with the same VC being used as the control actuator, and a linear ramp request, this time over 300–550 ms (shot 49068). This is plotted in figure 3 as well; the plasma boundary match to reference at 600 ms is $\eta = 6.9 \times 10^{-3}$. It is remarkable that this can been achieved with a single VC—the standard approach to generating VCs requires linearising around a GSE solution, and therefore is only valid for small deviations in strike point location. This means achieving a full sweep in feedback requires chaining together multiple VCs and managing the handover from one to another; the SH-derived VC is altogether simpler to implement and equally effective.

An exceptional feature of VCs generated in this way is that they are plasma-agnostic: they are derived without any knowledge of the core plasma shape or current distribution, as machine geometry alone is sufficient to form the constraint matrix G and therefore δ***x***. This makes them very useful for real-time control applications. For example, one could create
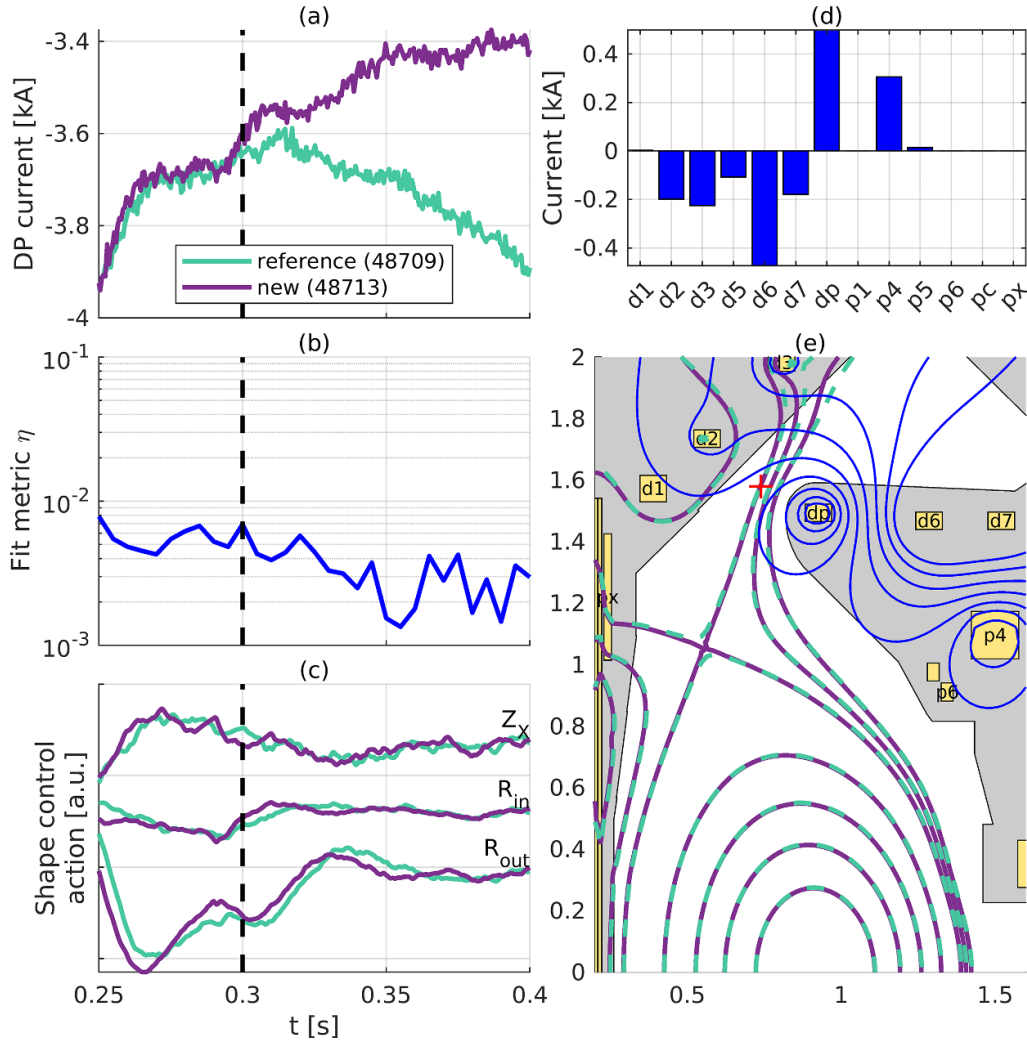
Figure 2. Experimental use of a VC derived using SH constraints to reduce DP current. (a) DP circuit current showing 500 A magnitude reduction at 400 ms (ramp in VC starts at dashed vertical line); (b) fit metric η (see (8)) comparing reference and new EFIT reconstructions; (c) control action applied to shape control VCs for both shots, with arbitrary *y*-axis offsets added; (d) PF current weightings for VC; (e) flux surfaces at 400 ms in reference (dashed) and new (solid) shots, with flux due to VC in blue and zero-flux constraint at location of red cross (+).

a standard, plasma-agnostic current limit avoidance VC for every PF circuit in the machine, to be applied in real-time during every shot. Similarly, the strike point control VC above can be expected to work for any core plasma (though controller gains may change), which avoids having to compute new VCs for each scenario or transferring VCs across scenarios.

3.3. Experimental scenario design

Having verified robust efficacy of the method in both simulation and experiment, we now show how SH constraints could be exploited to design ST divertor geometries for present-day experiments. We frame this task as a constrained optimisation problem for the PF currents to produce a certain core plasma and divertor geometry simultaneously; this is analogous to an inverse-GSE problem, though solving the GSE is now no longer a part of the optimisation. Instead, we have the

constraints (6) with general solution (7). Assuming a quadratic objective $||\mathsf{A}\boldsymbol{x} - \boldsymbol{b}||_2^2$, the optimum sits at

$$\boldsymbol{x} = \mathsf{N}(\mathsf{A}\mathsf{N})^+ \boldsymbol{b} + \left(\mathsf{I} - \mathsf{N}(\mathsf{A}\mathsf{N})^+ \mathsf{A}\right)\mathsf{G}^+ \boldsymbol{c}. \qquad (10)$$

This ignores bounds on $\boldsymbol{x}$ (i.e. current limits); to incorporate them, we find values of $\boldsymbol{x}$ which are beyond limits, replace them with constraints at the limit, then re-evaluate (10).

In figure 4 we give numerical examples of solving this problem using MAST-U shot 45272—a 750 kA ELMy H-mode with conventional divertor leg—at 750 ms as a reference case. Our constraint terms G and $\boldsymbol{c}$ now consist of rows which enforce (i) that the first 4 odd SHs match those of the reference case; (ii) various total flux and field constraints in the divertor chamber—see the caption of figure 4; (iii) that P1/PC/PX/P6 currents are unchanged. We choose Ohmic dissipation in the
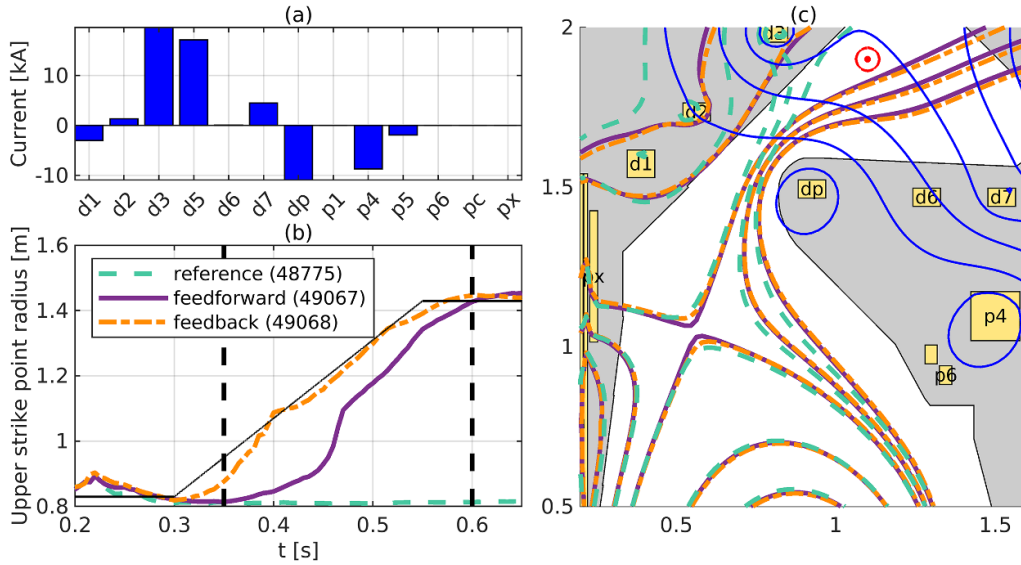
Figure 3. Decoupled divertor control in feed-forward (solid purple) and feedback (dash-dotted orange) compared to reference (dashed turquoise). (a) PF current weightings for VC; (b) evolution of strike point radius—note that feed-forward ramp (between vertical dashed black lines) is delayed by 50 ms with respect to the feedback request (thin solid black line); (c) flux surfaces at 600 ms for all three shots, with flux surfaces due to VC in blue and flux control point as a red target.

PF circuits as the minimisation objective and enforce MAST-U control system current limits.

The free-boundary equilibrium is recomputed after the optimisation of the PF currents to validate that η stays below our heuristic bound of 10^{-2}, and that divertor constraints are obeyed. Figure 4 shows the generation of a wide range of divertor geometries, which are decoupled from the core plasma shape. The key advantages of this approach are:

- Intuitive: rather than requiring the user to define multiple boundary control points and weights, a single input—the number of SHs constrained—determines the quality of core fit. A palette of exact constraints provide a natural way to create the desired divertor geometry;
- Rapid: the problem is a simple constrained least squares optimisation and its solution (10) amounts to multiplication by a matrix with fewer elements than the number of PF coils squared, i.e. instantaneous on a human timescale and an excellent candidate for real-time application;
- Robust: as shown in figure 1(a) good core fit is practically guaranteed for moderate $\ell_{\max}$, yet avoiding iteration of the non-linear GSE means there are no issues with convergence.

There are some limitations, however:

- No flexibility to change the core plasma shape—a separate inverse or forward GSE solve is required to produce the base case;
- In its simplest form, no freedom to adjust centre-column PF currents, though this is addressed in the appendix;
- Limited utility for modifying inboard divertor legs or the magnetic field around the X-point (e.g. to create a snowflake

divertor [32, 33]), as they sit within the imaginary sphere containing the plasma.

3.4. Coil position optimisation

Our final example considers the natural extension of this optimisation problem to one in which the coil positions, as well as currents, are free to vary. Such an exercise is an essential part of reactor design, in which the need to realise a particular plasma scenario (core and exhaust) must be balanced against the magnetic cage cost (number and sizes of coils) and spatial integration challenges. Existing tools to solve this problem [34, 35] involve converging the GSE concurrently with the non-linear optimisation, and tend to require user effort and expertise to formulate, with the calculation itself often suffering from instability, a lack of robustness, and fairly long run times. We propose a two-step approach which alleviates these issues:

(i) An inverse-GSE solve is performed with fixed PF coil positions, to create a core plasma with the required shape and internal profiles. The coilset here is arbitrary, as we simply use it to extract the SHs required;

(ii) Formulate an optimisation problem for the PF coil positions and currents which satisfy a set of linear and non-linear equality and inequality constraints, which minimally include the requirement to match the first few SHs of the solution to (i).

Step (i) is routine for any GSE solver, and step (ii) is straightforward because all constraints are closed-form expressions. In particular, the SH constraints are given by (4), with the filament positions r_f, θ_f (equivalently R_f, Z_f) as well
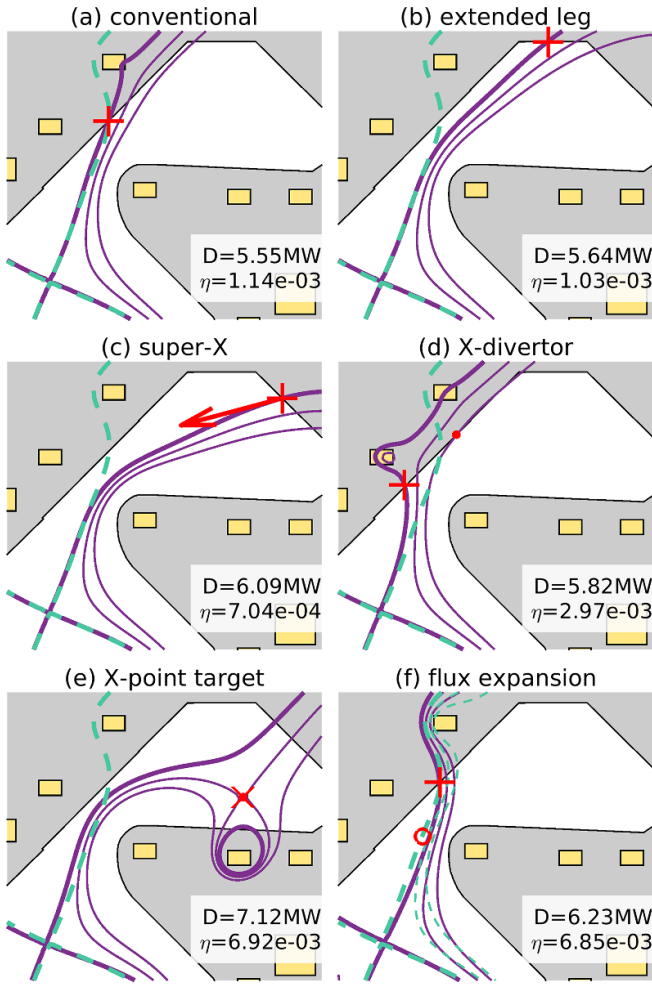
Figure 4. Numerical examples of divertor geometry design and optimisation. Dashed turquoise lines are the base case and solid purple lines the optimised equilibria; thick lines are the plasma separatrix and thin lines at normalised flux of 1.05 and 1.1. Red symbols denote divertor constraints: crosses ($+$) are separatrix constraints, the arrow in (c) is a poloidal field direction constraint, dots are normalised flux constraints (1.1 in (d), 1.05 in (e)), the $\times$ in (e) is an X-point constraint, and the circle in (f) is a poloidal field constraint to 75% of the base case value. The Ohmic dissipation D (which is 6.39 MW in the base case) and core plasma fit metric η are shown.

as I_f being free variables, and the coefficients A^ℓ_{PF} known from (i).

To demonstrate this process we manufacture a toy problem based around MAST-U. Specifically, we ask: can we take a 1MA ST plasma (shot 48281 at 400 ms) and design a coilset to reproduce it, with an extended divertor leg, and using far fewer PF coils? In essence, the equilibrium reconstruction is taking the place of (i), and we focus on solution of (ii). Our equality constraints are (i) the first 4 odd SHs, (ii) a flux constraint at the desired strike point, and (iii) explicit up-down symmetry. Inequality constraints are used to create bounding boxes for each coil and enforce current limits. We choose to minimise the sum of the currents squared, and use MATLAB's `fmincon` optimisation function with the 'active-set' algorithm to solve
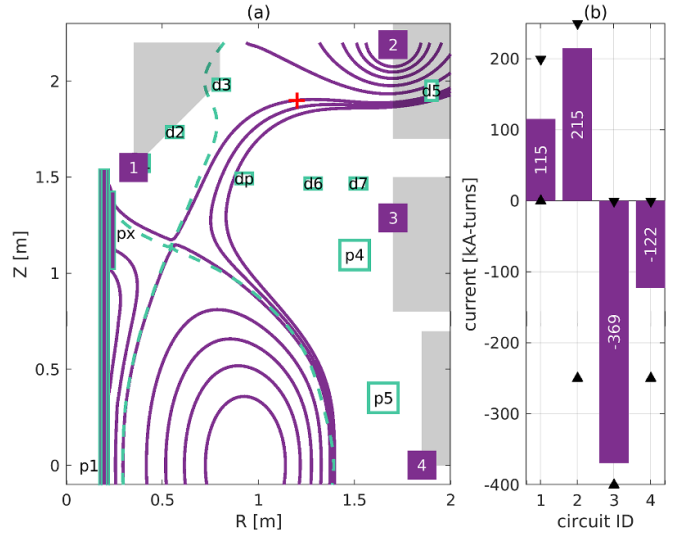

Figure 5. PF coil position optimisation. (a) Upper poloidal cross-section: dashed turquoise line is original plasma boundary, unfilled turquoise rectangles are original coilset, red $+$ is strike-point constraint, grey areas are go-zones for each coil, filled purple rectangle are optimised coil locations, solid purple lines are re-calculated equilibrium flux surfaces; (b) optimised PF currents and limits (black triangles).

the problem. Figure 5 illustrates both the problem formulation and the converged optimum.

There are a few interesting points to note:

- The inboard coils (including currents) are unchanged from the base case. In a reactor design exercise, the inboard would have to be designed as a precursor to stage (i) above (though see appendix);
- The optimiser has chosen to place coil 4, and therefore its up-down symmetric partner, at $Z = 0$—the two have become one midplane coil carrying -244 kA-turns. Hence the solution has 10 coils, compared to 21 for the real tokamak;
- The core plasma is not as well-matched as previous examples ($\eta = 1.15 \times 10^{-2}$ here) because with fewer coils further from the plasma, the higher (unconstrained) SHs are inevitably smaller and therefore further from the base case values;
- The optimisation run-time (which does not include GSE solves for initialisation and verification) is less than 1 s on an Intel Xeon E5-2640v3 CPU, and could easily be accelerated if required by applying a gradient-based method as all constraints are differentiable.

Although this example is rather artificial, it demonstrates a very powerful workflow for solving the coil position optimisation problem which shares benefits with the scenario design problem above: it requires minimal expertise to formulate, converges rapidly, and is highly robust.

4. Discussion and conclusions

The three goals identified in our introduction as essential for a performant divertor magnetic geometry in ST reactors—real-time control, experimental tests, and coilset design—have a loose 1:1 mapping to the three main advantages of the SH method—speed, user-friendliness, and robustness. We now discuss how each can benefit from application of the SH methods outlined here, and opportunities for future exploitation.

4.1. Real-time control

Because the SH approach does not rely on perturbation to or linearisation around a GSE solution [36, 37], it produces plasma-agnostic divertor control VCs which work across a wide range of exhaust magnetic configurations, and are quick and easy to implement in real-time. This gives opportunities experimentally and in future reactors for effective exhaust heat load mitigation including control of the strike point location (either static or sweeping), flux expansion (e.g. for detachment control [38]), coil current limit avoidance, and many others. An interesting possibility is to create an optimal, plasma-agnostic strike point control VC for all possible strike points and put them in a look-up table for the real-time controller to dynamically change its actuator throughout a shot.

Also in prospect are core plasma shape control improvements through application of SHs. We have learnt that high-order SHs have little bearing on the core plasma shape, so it makes sense to use only lower-order ones when controlling shape parameters such as inner/outer radii and X-point location. This provides a means to regularise the shape controller so it does not waste effort changing high-order SHs, and therefore operates as efficiently as possible. A similar approach applies for cancelling out the solenoid stray field for real-time current control (see appendix).

4.2. Experimental tests

The intuitive, user-friendly approach to designing ST divertor geometries demonstrated in section 3.3 lends itself to inter-shot use in present-day experiments looking at the relative merits of candidate advanced exhaust solutions in a way which is decoupled from the core plasma. The least squares optimisation formulation has a number of advantages beyond speed:

- Exact divertor constraints allow systematic variation of some magnetic geometry parameters whilst fixing others;
- There is flexibility in the choice of objective function (entirely distinct from plasma shape constraints) depending on user requirements;
- Enforcing bounds is trivial, reducing the chance of coil current limit violations;
- With a relatively simple optimisation problem, it is a more trustworthy result when the optimiser reports that the problem is infeasible.

Taken together, these make SH constraints a powerful tool to aid future ST divertor experiments by greatly reducing the number of shots required for scenario development.

4.3. Reactor design

The task of optimising PF coil positions and currents alongside the plasma shape and divertor solution for reactor-scale tokamaks contains multiple sources of non-linearity, making it challenging to converge and sensitive to input parameters which must be expertly chosen by the user. Applying SH constraints avoids the greatest source of this problem—the GSE solve—and therefore, by splitting the problem into two stages, increases the robustness with which the global optimum is found. This constitutes an important step forward in coilset design algorithms and will enable great cost reductions through the minimisation of number and volume of PF magnets, and accelerating the resolution of spatial integration issues; work is ongoing to incorporate SH constraints within the BLUEMIRA reactor design framework [39].

In future reactors, as to a lesser extent for present-day experiments, it will be important to compensate for currents induced in passive conductors during dynamic phases of plasma discharges. Here again SHs may help, since they can be used to convert the non-linear 'inverse evolution' problem [40] into a requirement on the vacuum field evolution which may be solved in much the same way as [41].

Finally, we highlight that extension of all the ideas within this work to conventional aspect ratio tokamaks will be possible by applying toroidal harmonics [42] in place of SHs.

Data availability statement

Acknowledgments

Appendix. Centre-column circuit compensation

As highlighted in section 2, the expression (4) cannot be directly applied to PF coils within an imaginary sphere which just encloses the core plasma, i.e. the central solenoid and any inboard divertor and X-point control coils. Nevertheless, it is possible to calculate a set of SH coefficients which, when added to the field due to a given centre-column circuit, result
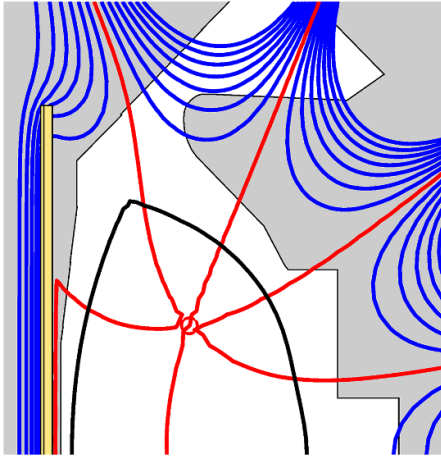
Figure A1. Contours of poloidal flux (blue, with hexapole null highlighted in red) due to the MAST-U central solenoid (yellow), plus the four odd SHs which give a total $\boldsymbol{B}_p = \nabla \boldsymbol{B}_p = \boldsymbol{0}$ at the red circle. The boundary of shot 45272 at 750 ms is shown in black.

in a very small poloidal field—though a finite poloidal flux—over the core plasma region. This is much the same as the common practice in existing tokamaks of using combinations of external PF currents to compensate for the stray field due to the solenoid, except that now the field required is defined in terms of SHs rather than circuit currents. In figure A1 we show this idea applied to the MAST-U central solenoid; we work out the first four odd SH coefficients such that, when added to the solenoid field, they give zero poloidal field and field gradients (a hexapole null) at a point in the centre of each half-plane.

From a real-time control point of view, a SH-based controller could use rate of change of the field in figure A1 to control the plasma current, with the request on the SHs required being passed to another controller which appropriately distributes current amongst the external PF circuits; this gives more flexibility than having a dedicated set of coils or VC for this purpose.

From a scenario design point of view, the negation of the coefficients calculated here can be used for the centre-column circuit in place of the regular SH coefficients evaluated via (4). This works because the poloidal *field* over the core plasma is broadly the same as that due to the centre-column circuit, though care must be taken because the poloidal *flux* will change, with consequences for inductively-driven plasmas.

ORCID iDs

O P Bardsley https://orcid.org/0000-0002-1525-675X
J L Baker https://orcid.org/0000-0002-5580-2073
C Vincent https://orcid.org/0000-0002-7227-409X

References

[1] Windridge M 2019 Smaller and quicker with spherical tokamaks and high-temperature superconductors *Phil. Trans. R. Soc.* A **377** 20170438

[2] Menard J E *et al* 2016 Fusion nuclear science facilities and pilot plants based on the spherical tokamak *Nucl. Fusion* **56** 106023

[3] Wilson H, Chapman I, Denton T, Morris W, Patel B, Voss G and Waldon C (the STEP Team) 2020 STEP — on the pathway to fusion commercialization *Commercialising Fusion Energy: How Small Businesses are Transforming Big Science* (IOP Publishing)

[4] Morris W, Harrison J R, Kirk A, Lipschultz B, Militello F, Moulton D and Walkden N R 2018 MAST Upgrade divertor facility: a test bed for novel divertor solutions *IEEE Trans. Plasma Sci.* **46** 1217–26

[5] Kurskiev G S *et al* 2021 Energy confinement in the spherical tokamak Globus-M2 with a toroidal magnetic field reaching 0.8 T *Nucl. Fusion* **62** 016011

[6] Creely A J *et al* (the SPARC Team) 2020 Overview of the SPARC tokamak *J. Plasma Phys.* **86** 865860502

[7] Buxton P F, Connor J W, Costley A E, Gryaznevich M P and McNamara S 2019 On the energy confinement time in spherical tokamaks: implications for the design of pilot plants and fusion reactors *Plasma Phys. Control. Fusion* **61** 035006

[8] Bruzzone P, Fietz W H, Minervini J V, Novikov M, Yanagi N, Zhai Y and Zheng J 2018 High temperature superconductors for fusion magnets *Nucl. Fusion* **58** 103001

[9] Sykes A *et al* 2018 Compact fusion energy based on the spherical tokamak *Nucl. Fusion* **58** 016039

[10] Coleman M and McIntosh S 2019 BLUEPRINT: a novel approach to fusion reactor design *Fusion Eng. Des.* **139** 26–38

[11] Faugeras B 2020 An overview of the numerical methods for tokamak plasma equilibrium computation implemented in the NICE code *Fusion Eng. Des.* **160** 112020

[12] Hofmann F 1988 FBT - a free-boundary tokamak equilibrium code for highly elongated and shaped plasmas *Comput. Phys. Commun.* **48** 207–21

[13] Welander A, Olofsson E, Sammuli B, Walker M L and Xiao B 2019 Closed-loop simulation with Grad-Shafranov equilibrium evolution for plasma control system development *Fusion Eng. Des.* **146** 2361–5

[14] Ambrosino G *et al* 2005 XSC plasma control: tool development for the session leader *Fusion Eng. Des.* **74** 521–5

[15] Khayrutdinov R R and Lukash V E 1993 Studies of plasma equilibrium and transport in a tokamak fusion device with the inverse-variable technique *J. Comput. Phys.* **109** 193–201

[16] Bardsley O P and Hender T C 2020 On the axisymmetric stability of tokamaks with ferromagnetic walls *Phys. Plasmas* **27** 102508

[17] Volegov P L, Danly C R, Fittinghoff D, Geppert-Kleinrath V, Grim G, Merrill F E and Wilde C H 2017 Three-dimensional reconstruction of neutron, gamma-ray and x-ray sources using spherical harmonic decomposition *J. Appl. Phys.* **122** 175901

[18] Li H, Liang G and Huang Y 2021 An efficient radiation analysis approach through compressive model for laser driven inertial confinement fusion *Comput. Phys. Commun.* **259** 107644

[19] Li F X, Voccio J P, Ahn M C, Hahn S, Bascuñán J and Iwasa Y 2015 An analytical approach towards passive ferromagnetic shimming design for a high-resolution NMR magnet *Supercond. Sci. Technol.* **28** 075006

[20] Thébault E, Hulot G, Langlais B and Vigneron P 2021 A spherical harmonic model of Earth's lithospheric magnetic field up to degree 1050 *Geophys. Res. Lett.* **48** e2021GL095147

[21] Finlay C C, Kloss C, Olsen N, Hammer M D, Tøffner-Clausen L, Grayver A and Kuvshinov A 2020 The

CHAOS-7 geomagnetic field model and observed changes in the South Atlantic Anomaly *Earth Planets Space* **72** 156

[22] Jackson J D 1998 *Classical Electrodynamics* 3rd edn (Wiley)

[23] Lao L L, John H S, Stambaugh R D, Kellman A G and Pfeiffer W 1985 Reconstruction of current profile parameters and plasma shapes in tokamaks *Nucl. Fusion* **25** 1611

[24] Kogan L *et al* (MAST Upgrade team) 2022 First MAST-U equilibrium reconstructions using the EFIT++ code: 48th European physical society conference on plasma physics, EPS 2022 *48th EPS Conf. on Plasma Physics (27 June–1 July 2022)* (European Physical Society (EPS))

[25] Berkery J W, Sabbagh S A, Kogan L, Ryan D, Bialek J M, Jiang Y, Battaglia D J, Gibson S and Ham C 2021 Kinetic equilibrium reconstructions of plasmas in the MAST database and preparation for reconstruction of the first plasmas in MAST upgrade *Plasma Phys. Control. Fusion* **63** 055014

[26] Cunningham G 2013 High performance plasma vertical position control system for upgraded MAST *Fusion Eng. Des.* **88** 3238–47

[27] Bardsley O P 2023 Metric for goodness of fit of approximations of MAST-U plasma reconstructions using vacuum fields truncated at varying numbers of spherical harmonics (Version 1) United Kingdom Atomic Energy Authority (UKAEA) (https://doi.org/10.14468/P80T-J063)

[28] McArdle G, Pangione L and Kochan M 2020 The MAST Upgrade plasma control system *Fusion Eng. Des.* **159** 111764

[29] Anand H, Eldon D, Kochan M, McArdle G, Pangione L and Wang H Q 2022 Validation of the strike point position estimation with the local expansion method for MAST upgrade on the DIII-D tokamak *Fusion Eng. Des.* **177** 113086

[30] Kochan M, Anand H, Lvovskiy A, Ryan P, Verhaegh K, Wijkamp T, Kirk A and McArdle G (MAST Upgrade Team) 2023 Real-time plasma shape reconstruction on MAST Upgrade based on local expansion *Proc. 30th IEEE Symp. on Fusion Engineering*

[31] McArdle G *et al* (the MAST Upgrade Team) 2023 Integrated real-time control on the MAST Upgrade tokamak *29th IAEA Fusion Energy Conf. (FEC 2023) (London, 16–21 October 2023)*

[32] Ambrosino R, Albanese R, Coda S, Mattei M, Moret J M and Reimerdes H 2014 Optimization of experimental snowflake configurations on TCV *Nucl. Fusion* **54** 123008

[33] Albanese R, Ambrosino R and Mattei M 2014 A procedure for the design of snowflake magnetic configurations in tokamaks *Plasma Phys. Control. Fusion* **56** 035008

[34] Coleman M and McIntosh S 2020 The design and optimisation of tokamak poloidal field systems in the BLUEPRINT framework *Fusion Eng. Des.* **154** 111544

[35] Franza F, Boccaccini L V, Fable E, Landman I, Maione I A, Petschanyi S, Stieglitz R and Zohm H 2022 MIRA: a multi-physics approach to designing a fusion power plant *Nucl. Fusion* **62** 076042

[36] Albanese R and Villone F 1998 The linearized CREATE-L plasma response model for the control of current, position and shape in tokamaks *Nucl. Fusion* **38** 723

[37] Heumann H, Blum J, Boulbe C, Faugeras B, Selig G, Ané J M, Brémond S, Grandgirard V, Hertout P and Nardon E 2015 Quasi-static free-boundary equilibrium of toroidal plasma with CEDRES++: computational methods and applications *J. Plasma Phys.* **81** 905810301

[38] Anand H, Humphreys D, Eldon D, Leonard A, Hyatt A, Sammuli B and Welander A 2020 Plasma flux expansion control on the DIII-D tokamak *Plasma Phys. Control. Fusion* **63** 015006

[39] Graham G, Bardsley O, Khan S, Pearce A and Johnson M 2023 Conceptual design of spherical tokamaks using spherical harmonic flux constraints in Bluemira *Proc. 30th IEEE Symp. on Fusion Engineering*

[40] Blum J, Heumann H, Nardon E and Song X 2019 Automating the design of tokamak experiment scenarios *J. Comput. Phys.* **394** 594–614

[41] di Grazia L E and Mattei M 2022 A numerical tool to optimize voltage waveforms for plasma breakdown and early ramp-up in the presence of constraints *Fusion Eng. Des.* **176** 113027

[42] Alladio F and Crisanti F 1986 Analysis of MHD equilibria by toroidal multipolar expansions *Nucl. Fusion* **26** 1143