J. Plasma Phys. (2025), *vol*. 91, E151 © The Author(s), 2026. Published by Cambridge University Press.
This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited. doi:10.1017/S0022377825100925

Omnigenous umbilic stellarators

R. Gaur[1], D.G. Panici[1], T.M. Elder[2,3], M. Landreman[3], K. Unalmis[1], Y.G. Elmacioglu[1], D.W. Dudt[4], R. Conlin[3] and E. Kolemen[1,5]

[1]Princeton University, Princeton, NJ 08540, USA
[2]Max Planck Institute for Plasma Physics, 17491 Greifswald, Germany
[3]Institute for Research in Electronics and Applied Physics, University of Maryland, College Park, MD 20740, USA
[4]Princeton Plasma Physics Laboratory, Princeton, NJ 08540, USA
[5]Thea Energy, Princeton, NJ, USA

Corresponding author: Rahul Gaur, rgaur@terpmail.umd.edu, rg6256@princeton.edu

(Received 7 May 2025; revision received 8 September 2025; accepted 23 September 2025)

To better understand the dependence of the magnetic field structure in the plasma edge on the plasma boundary shape, in the context of X-point and island divertor designs, we define and develop a class of stellarators called umbilic stellarators. These equilibria are characterised by a single continuous high-curvature edge on the plasma boundary that goes around multiple times toroidally before meeting itself. We develop a technique that allows us to simultaneously optimise the plasma boundary along with a curve lying on the boundary on which we impose a high curvature while imposing omnigenity – a property of the magnetic field that ensures trapped particle confinement throughout the plasma volume. We find that umbilic stellarators naturally tend to favour piecewise omnigenity instead of omnigenity with a specific helicity. After generating omnigenous umbilic stellarators, we design coil sets for some of them and explore the fieldline structure in the edge and its sensitivity to small fluctuations in the plasma. Finally, using single-stage optimisation, we simultaneously modify the plasma and coil shape and propose an experiment to modify an existing tokamak to a finite-β stellarator using this technique and explore a potentially simpler way to convert a limited tokamak into a diverted stellarator.

Key words: fusion plasma, plasma confinement

1. Introduction

Stellarators (Spitzer 1958) are toroidal nuclear fusion devices that are used to study magnetic confinement fusion. Unlike axisymmetric tokamaks, stellarators possess discrete toroidal symmetry. The lack of continuous toroidal symmetry offers considerably greater design freedom, at the expense of trapped particle confinement, a quality guaranteed in axisymmetry through Noether's theorem (Noether

1918). Therefore, we seek stellarators with trapped particle confinement, a property known as omnigenity.

Omnigenous magnetic fields ensure zero bounce-averaged radial drift (Hall & McNamara 1975; Cary & Shasharina 1997), of which both axisymmetry and quasisymmetry are subsets (Boozer 1983; Nührenberg & Zille 1988; Dudt *et al*. 2024). Recent advances in stellarator optimisation have generated omnigenous configurations with excellent neoclassical confinement levels (Landreman & Paul 2022; Goodman *et al*. 2023). However, these configurations are typically only calculated up to the last closed flux surface (LCFS). To develop a stellarator fusion power plant, it is essential to also predict and optimise the magnetic field structure outside the LCFS (Parra Diaz *et al*. 2024). This is necessary to separate the high-temperature plasma from the plasma-facing components while efficiently disposing of fusion by-products like helium ash and impurities such as heavier ions.

The structure outside the plasma boundary depends on the boundary shape, so changing the boundary shape is a potential way to obtain a desirable fieldline structure in the edge. Indeed, divertors (Burnett *et al*. 1958) have been greatly successful in tokamaks and offer significantly improved performance over limiters (Wagner *et al*. 1982). Depending on the boundary shape and coil positions, one can design various divertor configurations in tokamaks such as the single-null, double-null, snowflake (Ryutov 2007) and super-X (Valanju *et al*. 2009). These configurations utilise poloidal and divertor coils to create and shift an X-point – a region in tokamaks where the poloidal field disappears.

Stellarators have utilised three distinct divertor configurations (Bader *et al*. 2018*a*): helical, resonant and non-resonant divertors. Helical divertors were used in heliotron/torsatron concepts such as the Heliac device (Harris *et al*. 1985) and LHD (Ohyabu *et al*. 1992). These divertors work by creating a stochastic region with edge surface layers and X-points, which the magnetic fieldlines cross before reaching the scrape-off layer. The coil design and magnetic configuration in the edge surface layers are complex and difficult to predict (Ohyabu *et al*. 1994). Moreover, applicability of this concept to omnigenous stellarators with non-helical symmetry still remains uncertain (Bader *et al*. 2018*a*), posing challenges for their design and optimisation.

Resonant divertors, used in W7-AS (Grigull *et al*. 2001; Feng *et al*. 2002) and W7-X (Pedersen *et al*. 2019), work by leveraging island formation and placing divertors between the X- and O-points of the islands. The island structure depends on the pitch of the fieldline in the boundary, whereas the connection length depends directly on the number of X-points and inversely on the magnetic fieldline pitch and the magnetic shear (Feng *et al*. 2011) making them useful for low-magnetic-shear stellarators.

Non-resonant divertors use a more complicated design to create a stochastic field in the edge, but similar to resonant divertors, they can be used to spread the heat load and improve plasma exhaust rejection. These divertors are also more resilient to changes in the magnetic field structure (Bader *et al*. 2017), but the magnetic field structure in the chaotic field region is complicated to calculate and requires further understanding. Recent work (Boozer & Punjabi 2018; Punjabi & Boozer 2022; Garcia *et al*. 2023; Davies *et al*. 2025) has further explored the possibility of using non-resonant divertors in stellarators.

Motivated by the need to understand how boundary shaping affects fieldline structure in the edge and X-point and island divertor design, in this paper, we study the umbilic stellarator (US) as a method of achieving control over the edge magnetic

https://doi.org/10.1017/S0022377825100925 Published online by Cambridge University Press

field topology, and generate USs for vacuum and finite-β cases.[1] Umbilic stellarators are characterised by regions of high curvature on the boundary and how these high-curvature regions are positioned. We develop a technique that allows us to obtain USs by simultaneously optimising a curve and the plasma boundary, and imposing a high curvature along this curve which creates a ridge along the plasma boundary. Additionally, we ensure that all equilibria have reasonable omnigenity, ensuring low neoclassical transport. Finally, we propose an experiment to modify an existing tokamak to a stellarator using our optimisation process and build it using umbilic coils, which are helically wound coils, similar to divertor coils, with fractional helicity that carry a small current (compared with the plasma current), and follow an umbilic edge on the plasma boundary. We show how an umbilic coil can create a high-curvature umbilic region in finite-pressure equilibria, acting as an alternative to modular coils, to create a ridge on the plasma boundary.

In § 2, we explain the ideal magnetohydrodynamic (MHD) equilibrium. In § 3, we introduce the umbilic boundary shape, how to parametrise them and how we approximate them with DESC, and how to solve the ideal MHD equations throughout the volume enclosed by the approximated shape. In § 4, we present a set of optimised vacuum ($\beta = 0$) and finite-β stellarators and analyse their properties such as omnigenity. To understand the magnetic field structure in the edge, we design modular coils for the vacuum equilibrium and trace the fieldlines throughout the volume. In order to assess the robustness of the edge magnetic field configuration, we introduce a dummy current source at the magnetic axis to emulate a plasma fluctuation. In § 5, we take the Columbia High-Beta Torus-Extended Pulse (HBT-EP) plasma shape and modify it into an US without changing the coilset used for the original tokamak, just by adding an umbilic coil. We study cases where the umbilic coil carries current in the same (co-I) and opposite (counter-I) direction as the plasma current I, study the effect of the umbilic current on edge iota and solve a single-stage problem for the co-I case, by simultaneously modifying the plasma boundary and umbilic coil shape. In § 6, we conclude our study by highlighting existing limitations and exploring potential avenues for overcoming them.

2. Ideal MHD equilibrium

In this section, we briefly explain how we define and calculate an ideal MHD equilibrium in a stellarator. We then explain how to locally vary the gradients of that equilibrium.

A divergence-free magnetic field $\boldsymbol{B}$ can be written in the Clebsch form (D'haeseleer *et al*. 2012):

$$\boldsymbol{B} = \nabla\psi \times \nabla\alpha. \tag{2.1}$$

We focus on solutions whose magnetic fieldlines lie on closed nested toroidal surfaces, known as flux surfaces. We label these surfaces using the enclosed toroidal flux ψ. On each flux surface, the lines of constant α coincide with the magnetic fieldlines. Thus, α is known as the fieldline label. We define $\alpha = \theta_{\text{PEST}} - \iota\zeta$, where θ_{PEST} is the PEST straight fieldline angle, ζ is the cylindrical toroidal angle and

$$\iota = \frac{\boldsymbol{B}\cdot\nabla\theta_{\text{PEST}}}{\boldsymbol{B}\cdot\nabla\zeta} \tag{2.2}$$

[1] Where $\beta = 2\mu_0 p/B^2$ is the ratio of the plasma to the magnetic pressure.

https://doi.org/10.1017/S0022377825100925 Published online by Cambridge University Press

is the pitch of the magnetic fieldlines on a flux surface, known as the rotational transform. Using the Clebsch form of the magnetic field, we solve the ideal MHD force balance equation:

$$\boldsymbol{j} \times \boldsymbol{B} = \nabla p, \tag{2.3}$$

where the plasma current $\boldsymbol{j} = (\nabla \times \boldsymbol{B})/\mu_0$ from Ampere's law, p is the plasma pressure and μ_0 is the vacuum magnetic permeability. Unlike an axisymmetric case, for stellarators we have to solve (2.3) as an optimisation problem. We achieve this with the DESC (Dudt & Kolemen 2020; Conlin *et al*. 2023; Dudt *et al*. 2023; Panici *et al*. 2023) stellarator optimisation suite through minimisation of $F \equiv \boldsymbol{j} \times \boldsymbol{B} - \nabla p$. DESC can simultaneously solve an equilibrium while optimising for multiple objectives such as MHD stability, quasisymmetry and more, without re-solving the equilibrium force balance equation (2.3) for each optimisation variable. To reduce the computational cost further, we also use the fact that stellarators possess a discrete toroidal symmetry such that every stellarator comprises several identical sections. The number of identical sections is denoted by the field period n_{FP}. Therefore, the force balance problem is only solved in a single field period. In the following sections, we utilise DESC to optimise for omnigenous USs.

3. Umbilic stellarators

Umbilic stellarators draw inspiration from an umbilic bracelet (or torus) (Zeeman 1976), a three-dimensional shape with a cross-section that is a deltoid curve, a three-cusped hypocycloid which rotates by $\Delta\theta = 2\pi/3$ poloidally after a single toroidal turn $\Delta\zeta = 2\pi$. The sharp edge lies exactly on the cusp. Therefore, the sharp edge has to go around three times toroidally before meeting itself. A stellarator generated using an umbilic boundary is potentially advantageous for divertor designs, offering a continuous divertor option with a long connection length. For a smooth magnetic field, the sharp edge of such a boundary would be impossible for a magnetic fieldline. Therefore, the fieldline pitch on the boundary matches the umbilic topology. That is, $\iota = n_{\mathrm{FP}} m/n = 1/3$, where $m, n \in \mathbb{Z}$ corresponding to the poloidal and toroidal motion of the fieldline and n_{FP}, the field period, is the measure of discrete symmetry of the stellarator. For a general umbilic shape, m/n can attain rational values besides $1/3$. Further details are provided in this section.

3.1. *Defining an umbilic surface and edge*

The initial shape of a stellarator symmetric umbilic-torus-like surface ∂V is parametrised as a rotating polygon-like with convex sides:

$$\begin{aligned} R_{\mathrm{US}} = R_0 + 2\varrho \cos\left(\frac{\pi}{2n}\right) \cos\Bigg(& \frac{\theta + \frac{\pi}{n}(2\lfloor n\theta/2\pi \rfloor + 1)}{2} \\ & + \frac{mn_{\mathrm{FP}}\zeta + r_1 \sin(n_{\mathrm{FP}}\zeta) + r_4 \sin(2n_{\mathrm{FP}}\zeta)}{n} \Bigg) \\ & - \varrho \cos\left(\frac{\pi}{n}\left(2\left\lfloor \frac{n\theta}{2\pi} \right\rfloor + 1\right) + \frac{mn_{\mathrm{FP}}\zeta + r_1 \sin(n_{\mathrm{FP}}\zeta) + r_4 \sin(2n_{\mathrm{FP}}\zeta)}{n}\right) \\ & + r_2 \cos(n_{\mathrm{FP}}\zeta), \end{aligned} \tag{3.1}$$

https://doi.org/10.1017/S0022377825100925 Published online by Cambridge University Press

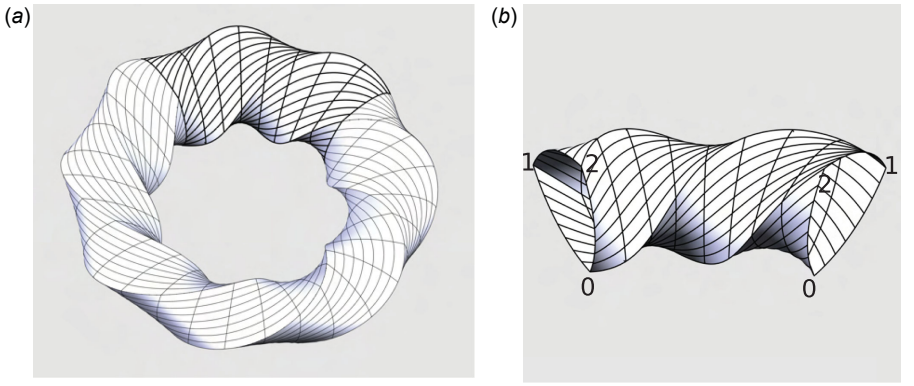


FIGURE 1. An umbilic boundary design with $n_{\mathrm{FP}} = 4$, $n = 3$, $m = 2$, $\varrho = 0.344$. All the rest of the parameters $r_j = 0$ in (3.1) and (3.2). In (*a*) we show the full shape and in (*b*) we present a single field period from the parametrisation. A point $p_i \in \{0, 1, 2\}$ corresponding to the sharp edge at each toroidal cross-section connects to $\mod(p_i + m, 3)$ after a single field period. The sharp edge defined by $R_{\mathrm{US}}(\theta = 0, \zeta)$, $Z_{\mathrm{US}}(\theta = 0, \zeta)$ meets itself after $n = 3$ toroidal turns.

$$Z_{\mathrm{US}} = 2\varrho \cos\left(\frac{\pi}{2n}\right) \sin\left(\frac{\theta + \frac{\pi}{n}(2\lfloor n\theta/2\pi \rfloor + 1)}{2} + \frac{mn_{\mathrm{FP}}\zeta + r_5 \sin(n_{\mathrm{FP}}\zeta)}{n}\right)$$
$$- \varrho \sin\left(\frac{\pi}{n}\left(2\left\lfloor \frac{n\theta}{2\pi} \right\rfloor + 1\right) + \frac{mn_{\mathrm{FP}}\zeta + r_5 \sin(n_{\mathrm{FP}}\zeta)}{n}\right) + r_3 \sin(n_{\mathrm{FP}}\zeta), \qquad (3.2)$$

where the major radius $R_0 = 1$ throughout this paper, unless stated otherwise, ϱ is the minor radius of the umbilic-torus-like surface, θ is the poloidal angle used in DESC, ζ is the cylindrical toroidal angle and $\lfloor \ldots \rfloor$ is the floor function, defined such that

$$\lfloor x \rfloor = \begin{cases} x, & \text{if } x \in \mathbb{Z} \\ N - 1, & \text{if } N - 1 < x < N \text{ and } x \notin \mathbb{Z}, \ N \in \mathbb{Z}. \end{cases} \qquad (3.3)$$

In the surface parametrisation, we also define n as the umbilic factor, which defines the number of sides in the analytical parametrisation, m is an integer that determines the poloidal motion of the sharp umbilic edge and n_{FP} is the field period. The periodicity of the edge is $n_{\mathrm{FP}}(m/n)$ – the edge closes on itself after $n/\gcd(n, n_{\mathrm{FP}})$ toroidal turns, where gcd is the greatest common divisor. The parameters r_1, r_4, r_5 define the convexity/concavity of a cross-section and how it changes toroidally, r_4 determines the magnetic axis torsion and r_2 is another parameter that determines the shape of the axis. The shape of such a boundary, the naming convention we use and periodicity of the edge are illustrated in figure 1.

The floor function creates the sharp umbilic edge, while the rest of the parameters grant us enough freedom that we can start with a boundary shape that has good quasisymmetry and allows us to create an umbilic edge with any low-order rational periodicity $n_{\mathrm{FP}}(m/n)$. It is also stellarator-symmetric, i.e. $R_{\mathrm{US}}(\theta, \zeta) = R_{\mathrm{US}}(-\theta, -\zeta)$, $Z_{\mathrm{US}}(\theta, \zeta) = -Z_{\mathrm{US}}(-\theta, -\zeta)$.

https://doi.org/10.1017/S0022377825100925 Published online by Cambridge University Press

We also define a separate grid on the boundary $\rho = 1$, the umbilic edge, on which a sharp curvature is imposed, defined by values $(\hat{\theta}, \zeta)$ so that

$$\hat{\theta} = \frac{mn_{\mathrm{FP}}\zeta + \sum_k^{N_u} a_k \sin(n_{\mathrm{FP}}\zeta)}{n}, \quad \gcd(m, n) = 1. \tag{3.4}$$

For the surface parametrisation, $\theta = 0$ corresponds to the umbilic edge – initial $\hat{\theta}$ is calculated along this edge but it can change freely during optimisation as the optimiser varies the coefficients a_k. Note that θ is the DESC poloidal angle throughout this paper, unless stated otherwise.

However, creating and solving with a perfectly sharp edge is difficult with a fully spectral solver. Therefore, in the next section, we explain how we approximate umbilic boundary shapes with the DESC code.

3.2. *Approximating umbilic shapes with DESC*

DESC uses a Fourier–Zernike spectral representation to solve the force balance equation in the (ρ, θ, ζ) coordinate system, where $\rho = \sqrt{\psi/\psi_b}$, $\theta = \theta_{\mathrm{PEST}} - \Lambda(\rho, \theta, \zeta)$, $\zeta = \phi$, the cylindrical toroidal angle. The problem is defined in a cylindrical coordinate system (R, ζ, Z) by decomposing it into Fourier–Zernike spectral bases (see Panici *et al*. 2023 for more details):

$$R(\rho, \theta, \zeta) = \sum_{m=-M,n=-N,l=0}^{M,N,L} R_{lmn} \mathcal{Z}_l^m(\rho, \theta) \hat{F}^n(\zeta), \tag{3.5}$$

$$Z(\rho, \theta, \zeta) = \sum_{m=-M,n=-N,l=0}^{M,N,L} Z_{lmn} \mathcal{Z}_l^m(\rho, \theta) \hat{F}^n(\zeta), \tag{3.6}$$

where l, m and n are the radial, poloidal and toroidal mode numbers, whereas L, M and N define the largest values of l, m and n, respectively, and $\mathcal{Z}$, $\hat{F}$ are Zernike and Fourier basis functions. These parameters define the resolution of a DESC equilibrium.

To solve the ideal MHD equation inside the umbilic shape, we have to first fit the Fourier–Zernike basis at $\rho = 1$ to the umbilic boundary, which requires solving a set of linear equations,

$$\begin{bmatrix} \sum_{m=-M,n=-N,l=0}^{M,N,L} R_{lmn} \mathcal{Z}_l^m(\rho = 1, \theta) \hat{F}^n(\zeta) \\ \sum_{m=-M,n=-N,l=0}^{M,N,L} Z_{lmn} \mathcal{Z}_l^m(\rho = 1, \theta) \hat{F}^n(\zeta) \end{bmatrix} = \begin{bmatrix} R_{\mathrm{US}}(\theta, \zeta) \\ Z_{\mathrm{US}}(\theta, \zeta) \end{bmatrix}, \tag{3.7}$$

to find R_{lmn}, Z_{lmn} so that the boundary shape matches with the umbilic parametrisation. Similarly, to obtain a_k, we parametrise the umbilic edge using $R_{\mathrm{curve}} = R(\rho = 1, \theta = 0, \zeta)$, $Z_{\mathrm{curve}} = Z(\rho = 1, \theta = 0, \zeta)$, and perform an inverse Fourier transform to obtain a_k. Note that upon fitting a spectral basis, the umbilic edge is smoothed out. After representing the umbilic shape in the Fourier–Zernike basis, we solve the ideal MHD equation inside the volume. Figure 2 illustrates this process of fitting and solving force balance in a US131 boundary shape.

The sharpness of the umbilic edge increases as we increase the resolution of the fit surface and the agreement between DESC and the umbilic parametrisation improves. Alternatively, one may use a double-Chebyshev (Fortunato 2024) instead of a double-Fourier representation to improve the spectral fit, local basis methods such as finite

https://doi.org/10.1017/S0022377825100925 Published online by Cambridge University Press

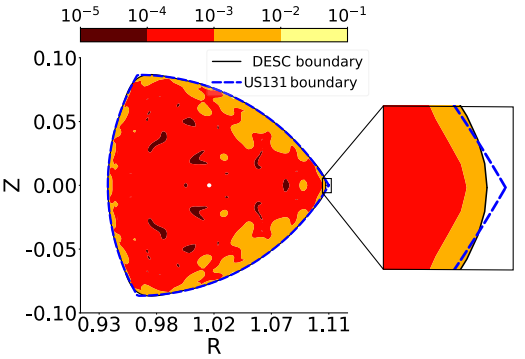


FIGURE 2. Contour plot of the flux surface average force balance error $|F|$ plotted on the cross-section $\zeta = 0$ of a vacuum DESC equilibrium created using a US131 boundary with $\rho = 0.1$, $r_1 = 0.133$, $r_5 = 0.1$. The actual boundary (in solid black) is obtained after fitting a Fourier–Zernike series to the true umbilic boundary (in dashed blue). The resolution of the DESC equilibrium is $L = M = N = 14$ and the resolution of the umbilic curve is $N_u = 12$.

element to represent and solve the equilibrium or use a nonlinear transformation to map the circular torus in the Fourier–Zernike basis to an umbilic shape, similar to the transformation used for circular tokamaks by Webster & Gimblett (2009). However, using a new basis and performing equilibrium optimisation are beyond the scope of this work.

In the following section, we discuss how, beginning with the parametrisation of the umbilic surface, we can optimise for omnigenous USs.

3.3. *Optimising an umbilic equilibrium with DESC*

As we noted in the previous section and figure 2, after fitting a Fourier–Zernike basis to the umbilic boundary, it becomes smooth because approximating a sharp boundary with a double-Fourier representation is usually difficult. Therefore, we have to modify the optimisation process for an umbilic equilibrium to maintain a high curvature, instead of a perfectly sharp umbilic edge. Additionally, for these shapes to be practically feasible fusion reactors, we also need good omnigenity. Therefore, we use the omnigenity and effective ripple objectives to improve the omnigenity throughout the volume.

After fitting to a Fourier–Zernike basis, the high-curvature edge, in general, does not align with the fieldline. In the unoptimised shape defined in (3.1) and (3.2), the ι is typically too small because of weak shaping. Therefore, we have to optimise separately for a boundary rotational transform to be close $n_{\mathrm{FP}}m/n$. Also, to ensure that the umbilic edge curve matches with the fieldline labelled $\alpha = 0$, we enforce a strong negative curvature $\kappa_{2,\rho}$ along the umbilic edge defined by (3.4). The definitions of the principal curvatures are given in Appendix A. The overall objective function is

$$\mathcal{F}_{\text{stage-one}} = w_A\, f^2_{\text{aspect}} + w_\iota\, f^2_{\text{iota}} + w_O\, f^2_{\text{om}} + w_U\, f^2_{\text{umbilic}} + w_R\, f^2_{\text{ripple}}, \tag{3.8}$$

where f_x are various objectives on the right-hand side for aspect ratio, boundary rotational transform, omnigenity, high curvature along the umbilic edge and effective ripple and w_A, w_ι, w_O, w_U, w_R are weights used with each objective function. The exact

https://doi.org/10.1017/S0022377825100925 Published online by Cambridge University Press

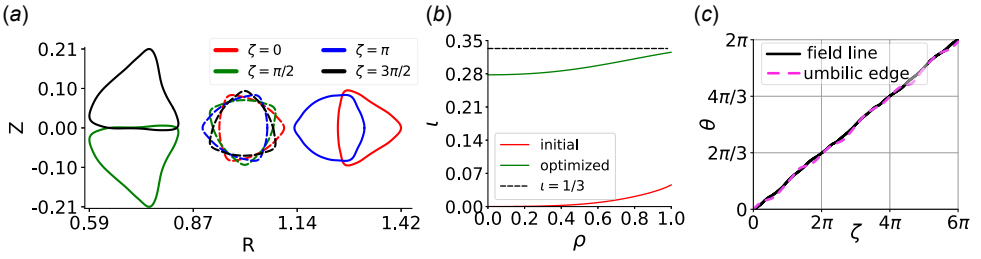


FIGURE 3. The US131 optimisation results showing a significant increase in magnetic axis torsion in (a) from the initial (dotted) to the optimised (solid) boundary. This leads to an increase in the rotational transform, as shown in (b) such that $\iota(\rho = 1)$ approaches $1/3$. The shaping also leads the fieldlines to align with the umbilic edge in the optimised equilibrium as shown in (c).

definitions of these objectives are provided in Appendix B. At each iteration of a DESC optimisation, $\mathcal{F}_{\text{stage-one}}$ is minimised while satisfying (2.3). Mathematically,

$$\min \mathcal{F}_{\text{stage-one}}(\hat{\boldsymbol{p}}), \qquad \text{s.t.} \quad \boldsymbol{\nabla}\left(\mu_0 p + \frac{B^2}{2}\right) - \boldsymbol{B}\cdot\boldsymbol{\nabla}\boldsymbol{B} = 0, \quad \psi_{\text{b}} = \psi_{\text{b0}}, \tag{3.9}$$

where ψ_{b0} is the user-provided enclosed toroidal flux by the boundary and $\hat{\boldsymbol{p}} \in \{R_{\text{b},mn}, Z_{\text{b},mn}, a_k\}$ comprises parameters that determine the boundary shape and the position of the umbilic curve on the boundary. With the objective $\mathcal{F}_{\text{stage-one}}$ defined, we run DESC on a single NVIDIA A100 GPU. A single optimisation takes less than 90 minutes.

4. Omnigenous umbilic stellarators

In this section, we develop two omnigenous umbilic equilibria using the process explained in the previous sections. The first equilibrium is a single-field-period vacuum stellarator with $n = 3$, $m = 1$ umbilic topology for which we also design coils, and the second equilibrium is a two-field-period finite-β stellarator with $n = 5$, $m = 2$ umbilic topology.

4.1. *A US131($n_{FP} = 1$, $n = 3$, $m = 1$) omnigenous vacuum stellarator*

We begin by calculating an omnigenous vacuum equilibrium with a single field period $n_{\text{FP}} = 1$, $n = 3$, $m = 1$. To do this, we start with the US131 boundary shape, using the parametrisation from § 3.1 with $\varrho = 0.1$ m, $r_1 = 0.133$, $r_5 = 0.1$ and the rest of the parameters $r_j = 0$. The toroidal flux enclosed by the initial equilibrium is $\psi_{\text{b}} = \pi\varrho^2$ (Tm^2). This shape is then approximated by DESC as explained in § 3.2 and optimised as described in § 3.3 with $w_{\text{O}} = 0$[2] to obtain an omnigenous equilibrium with an aspect ratio $A = 10.5$. The results are presented in figures 3 and 4.

To ensure that the high-curvature umbilic edge is aligned with the magnetic fieldlines, the boundary undergoes deformation, which results in an increased rotational

[2]By setting $w_O = 0$ and $w_R \neq 0$, we seek omnigenous configurations without constraining the magnetic field to a specific helicity. This improves the quality of the optimum. This could also be a reason we are more likely to get piecewise-omnigenous configurations.

https://doi.org/10.1017/S0022377825100925 Published online by Cambridge University Press

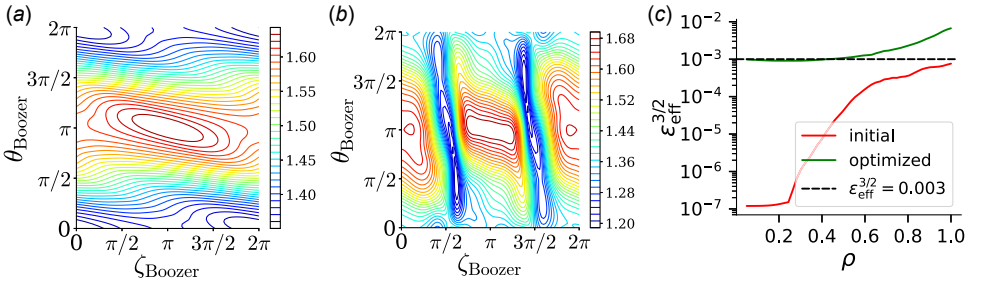


FIGURE 4. Output from the vacuum US131 optimisation showing Boozer plots on the boundary and the effective ripple. The output in (*b*) does not have a specific helicity but still achieves a low ripple transport.

transform. The optimised design minimises the objective $\mathcal{F}_{\text{stage-one}}$ while ensuring that the normalised second principal curvature is the lowest along the umbilic edge. The curvature on the umbilic edge is in the range $\kappa_{2,\rho} \in [-187, -70]\,\text{m}^{-1}$ whereas the average curvature on the rest of the boundary is $\text{avg}(\kappa_{2,\rho}) = -12\,\text{m}^{-1}$. This causes the magnetic fieldline to stay close to the umbilic edge throughout the range $\zeta \in [0, 6\pi]$ as shown in figure 3(*c*). In the limit where the edge is perfectly sharp, the fieldline will coincide with the umbilic edge.

The optimised US case is able to achieve an effective ripple ϵ_{eff} of less than 2 % ($\epsilon_{\text{eff}}^{3/2} < 0.0028$) for most of the volume ($\rho < 0.86$) and $\epsilon_{\text{eff}}^{3/2} \leqslant 0.001$ for $\rho \leqslant 0.6$. This ensures that neoclassical transport in the $1/\nu$ low-collisionality regime will not deteriorate the device performance in the core but may deteriorate confinement in the edge. An interesting observation is that we can achieve a low ripple transport despite the lack of a specific helicity of the B contours in the Boozer plot (figure 4*b*). This is an instance of a piecewise-omnigenous (Velasco *et al*. 2024) field that is far from quasisymmetry. Since we did not seek omnigenity with a specific helicity ($w_{\text{O}} = 0$), the effective ripple penalty acted as a proxy for omnigenity, leading to a piecewise-omnigenous configuration. To prove that these configurations are piecewise-omnigenous, we analyse the normalised second adiabatic invariant $\tilde{\mathcal{J}}_{\parallel}$ in Appendix C.

For a vacuum problem, obtaining the right fieldline pitch ($\iota = n_{\text{FP}} m/n$) on the boundary and ensuring force balance over a perfectly sharp (unapproximated) boundary make this problem difficult to solve compared with a finite-β equilibrium where one can specify the fieldline pitch throughout the volume. Therefore, we have compromised by creating a high-curvature edge instead of a perfectly sharp edge ($\kappa_{2,\rho} \to \infty$). If one were to use an alternative representation to represent a perfectly sharp boundary and solve this problem exactly, we would observe the rotational transform almost identical to the initial equilibrium in figure 4(*c*) sharply rising to $\iota = m\, n_{\text{FP}}/n$ at the edge.The behaviour of high-shear regions near sharp points would be similar to what is observed in tokamaks.

4.2. *Modular coil design for the US131 equilibrium*

In this section, we design modular coils for the US131 case using the DESC optimiser and elucidate the various difficulties faced when designing modular coils for

https://doi.org/10.1017/S0022377825100925 Published online by Cambridge University Press

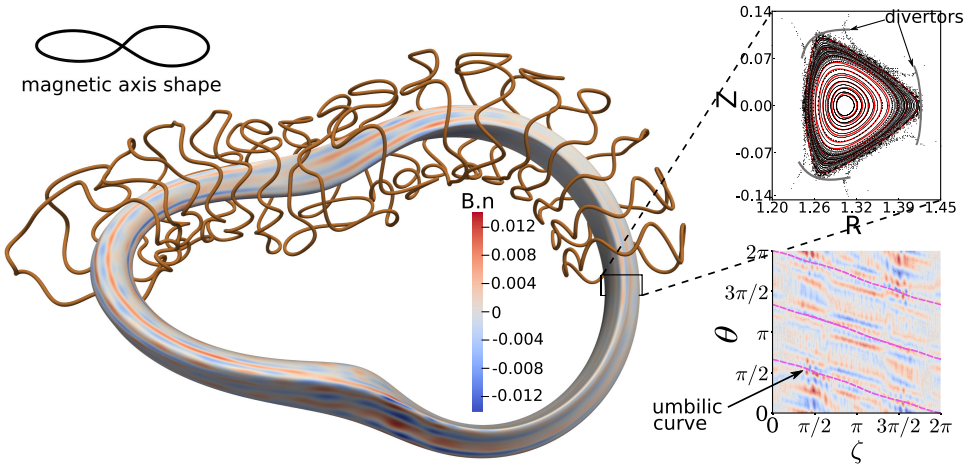


FIGURE 5. Coils and $\boldsymbol{B}_{\text{out}} \cdot \hat{\boldsymbol{n}}$ error plotted on the plasma boundary. Here $\max|(\boldsymbol{B}_{\text{out}} \cdot \hat{\boldsymbol{n}})/\boldsymbol{B}_{\text{out}}| \leqslant 1\,\%$. Only one group of unique coils are plotted. At the upper right we plot the Poincaré cross-section at $\zeta = 0$ along with the fixed-boundary solution (in red). The grey curves mark a favourable position to place a divertor. At the lower right we plot the $\boldsymbol{B}_{\text{out}} \cdot \hat{\boldsymbol{n}}$ error on a two-dimensional plane along with the umbilic edge (dashed purple).

strongly shaped USs. We create a coilset of 32 (stellarator-symmetric with 16 unique coils) coils using the Python implementation of `REGCOIL` (Landreman 2017; Panici *et al*. 2025) in `DESC`. Starting with this initial coilset, we then perform stage-two optimisation, where we fix the equilibrium boundary and allow the shape of the modular coils to change while minimising the objective

$$\mathcal{F}_{\text{stage-two}} = (B_{\text{out}}^2 - (B_{\text{in}}^2 + 2\mu_0 p))^2 + (\boldsymbol{B}_{\text{out}} \cdot \hat{\boldsymbol{n}})^2 \tag{4.1}$$

on the plasma boundary. Note that in stage-two optimisation, we ignore the surface current $\boldsymbol{K} = (\boldsymbol{B}_{\text{out}} - \boldsymbol{B}_{\text{in}}) \times \hat{\boldsymbol{n}}$. Therefore, to verify our solution, we calculate the fieldline trajectory due to the coils and create a Poincaré plot. The original fixed-boundary solution, its comparison with the Poincaré cross-section and the optimised coilset are presented in figure 5.

Note that even though the curvature along the umbilic ridge is high, it does not behave like a perfect X-point (or X-line) as the fieldline structure becomes chaotic in the edge and resembles the 'turnstiles' topology discussed by Punjabi & Boozer (2022). However, even though the fieldlines do not exit in the regions of the highest curvature, like a tokamak X-point, the Poincaré traces of the fieldlines stay aligned with the plasma boundary and 'flare' near the umbilic edge, which still makes the umbilic edge a favourable location to place divertors or limiters. This study indicates that the behaviour of magnetic fieldlines for a high-curvature ridge can differ significantly from that for a perfectly sharp edge.

Another observation is that modular coil design becomes complicated for umbilic equilibria as seen in the figure. The coils assume high curvature to reduce the normal field error along the umbilic edge. Increasing the number of coils improves the agreement of the fixed-boundary solution with the Poincaré cross-section and reduces the maximum coil curvature but also reduces the coil–coil distance. However, the number of modular coils needed for these umbilic equilibria already exceeds the number

https://doi.org/10.1017/S0022377825100925 Published online by Cambridge University Press

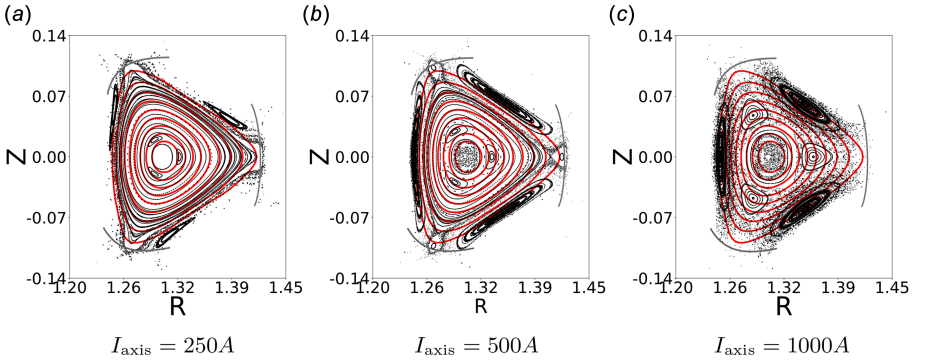


FIGURE 6. Poincare plots at $\zeta = 0$ with a current source placed on-axis for different values of the current, compared with the original fixed-boundary solution (in red). Increasing the current significantly reduces the nestedness of the flux surfaces but the fieldline structure in the edge still 'flares' out near the high-curvature edges and the divertor placement does not need to be modified.

of coils used for a single field period for a typical stellarator, which is almost always less than 10.

Finally, we note that the regions of maximum normal field error are parallel to the umbilic edge so that, on average, $\boldsymbol{b} \cdot \nabla(\boldsymbol{B}_{\text{out}} \cdot \hat{\boldsymbol{n}}) = 0$.[3] This pattern indicates a hidden relationship between the normal field error and the second principal curvature in USs.

To ensure practical coil design for USs, a better technique would be to perform single-stage optimisation, which simultaneously changes the plasma boundary and coil shapes, a union of $\mathcal{F}_{\text{stage-one}}$ and $\mathcal{F}_{\text{stage-two}}$, accounting for the coil complexity while maintaining omnigenity inside the plasma volume. This is done in § 5.2 for a simpler problem. Another possibility is to use umbilic coils, helically wound divertor-like coils with a fractional helicity (m/n) that carry a small current that pushes (pulls) the plasma away from (towards) itself, creating cusp-like regions, leading to a high-curvature umbilic edge. We do a similar optimisation to modify a tokamak in § 5.1, where we add an umbilic coil without changing the shape of the toroidal field coils to create an umbilic edge. In the next section, we use the same optimisation process to create a finite-β, omnigenous US.

4.3. *A US131 sensitivity analysis of the field structure in the edge*

In this section, we introduce a line current source along the magnetic axis of the US131 configuration. The line current aims to simulate an internal fluctuation within the plasma. For different values of the current source strength, we perform the fieldline integration across the entire volume. The purpose of this investigation is to assess the resilience of the fieldline structure at the edge of the US131 stellarator and to determine the extent of modification required for the divertor configuration. We present the fieldline structure for different values of the current in figure 6.

[3]We use the magnetic field unit vector instead of the vector along the umbilic edge because the magnetic field is reasonably aligned with the umbilic edge throughout its span as shown in figure 3(*c*).

https://doi.org/10.1017/S0022377825100925 Published online by Cambridge University Press

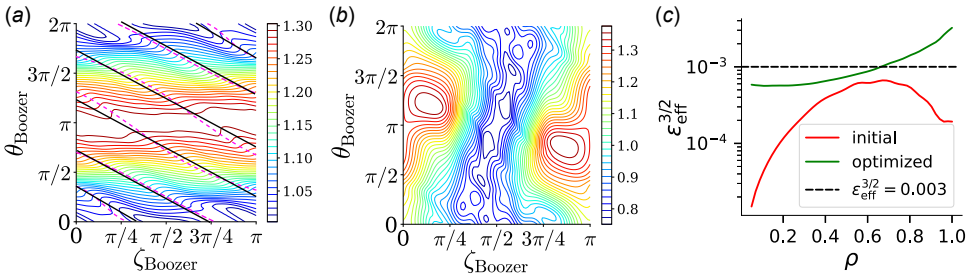


FIGURE 7. Optimisation output showing the (*a*) initial and (*b*) final magnetic field strength Boozer plots and (*c*) the effective ripple. The black lines in (*a*) correspond to the magnetic fieldlines and the dashed lines correspond to the umbilic curve. The optimised configuration is another instance of piecewise omnigenity.

We observe that even though the flux surfaces inside the plasma boundary lose nestedness and become stochastic, the divertor structure used for the unperturbed US131 case is still an effective choice for this equilibrium, i.e. the fieldline structure in the edge is resilient to fluctuations inside the plasma. Another observation is the emergence of 'O'-point-like structures near the high-curvature regions in figure 6(*b*).

4.4. *A US252* ($n_{FP} = 2, n = 5, m = 2$) *omnigenous finite-β stellarator*

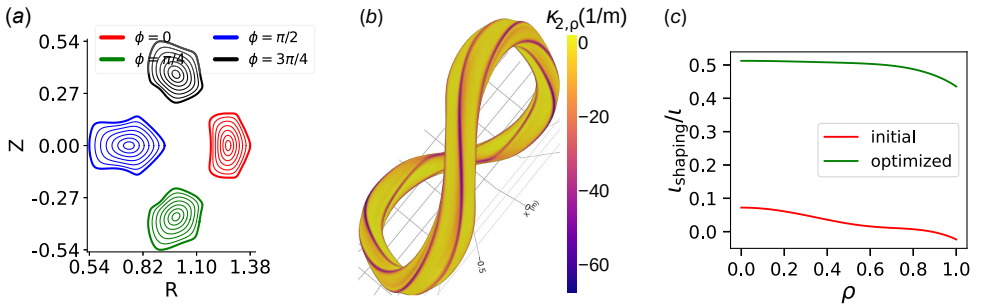
In this section, we generate a two-field-period $n = 5, m = 2$ finite-β stellarator. We start with the umbilic parametrisation defined in § 3.1 with $\varrho = 0.1$, $r_1 = 0.133$, $r_5 = 0.1$ and use a pressure profile $p = 5 \times 10^3(1 - \rho^2)$ Pa, a rotational transform $\iota = 0.755 + 0.043\rho^2$ and an enclosed toroidal flux $\psi_b = \pi\varrho^2 = 0.00314\,\text{T m}^2$. This creates an equilibrium with $\beta = 0.5\,\%$.

A finite-β stellarator provides more freedom compared with a vacuum case because the rotational transform can come from a combination of boundary shaping and plasma current. Therefore, we can achieve a higher rotational transform and use the additional freedom to achieve a much lower ripple transport compared with the vacuum case presented in the previous section. The optimisation uses the same objective function defined in (3.8) with the pressure and iota profiles fixed and $w_\iota = 0$. These results are presented in figure 7.

We can see in figure 7 that the ridges in the Boozer plots are only clearly visible in the initial equilibrium, which is close to quasisymmetry. The magnetic field of the optimised omnigenous equilibrium does not have any visible ridges in the Boozer plot, even though the curvature along the umbilic ridge increases after optimisation. This indicates that the appearance of ridges in the Boozer representation may only be a feature of quasisymmetric stellarators.

This equilibrium is optimised for poloidal omnigenity but, due to the umbilic curvature penalty, limits the poloidal omnigenity to the low-field region. The Boozer plot in figure 7(*b*) looks qualitatively like a piecewise-omnigenous Velasco *et al*. (2024) configuration, which looks similar to the magnetic field in the W7-X stellarator. The piecewise omnigenity of this equilibrium is demonstrated further in Appendix C. We also plot the cross-section of the equilibrium along with the second principal curvature in figure 8. The neoclassical transport coefficient $\epsilon_{\text{eff}}^{3/2} < 0.001$ for most ($\rho < 0.7$) of the volume, which means the equilibrium is omnigenous in the core.

https://doi.org/10.1017/S0022377825100925 Published online by Cambridge University Press



Optimized cross-sections Curvature on boundary ι fraction from shaping

FIGURE 8. Cross-section and three-dimensional shape of the optimised US252 equilibrium. The aspect ratio of the optimised equilibrium is $A = 6.5$. In (*c*), we see that around half of the rotational transform of the optimised equilibrium comes from shaping (Hirshman & Hogan 1986) reducing the dependence on external current drive.

We also observe that the optimised equilibrium deforms and rotates as we move toroidally, while maintaining high curvature along the umbilic curve. The curvature along the umbilic edge $\kappa_{2,\rho} \in [-67, -36]\,\mathrm{m}^{-1}$ while the average curvature on the rest of the plasma boundary is $\mathrm{avg}(\kappa_{2,\rho}) = -8\,\mathrm{m}^{-1}$. The variation in the curvature is broad, indicating that imposing high curvature at specific locations may be more suitable than others. This can be readily accomplished by selecting sections of the umbilic curve where high curvature is applied and could be a potential way to place potential locations for non-resonant divertors (Bader *et al*. 2018*b*).

5. Umbilic coil design for a finite-β equilibrium

As we have seen in § 4.1, creating modular coils for vacuum umbilic equilibria is a difficult task due to the sharp curvature regions. Therefore, in this section, we present an alternative strategy that uses umbilic coils to create a US111 finite-β equilibrium. Instead of modifying modular coils, we use an umbilic coil, a low-current, helically wound coil, similar to a divertor coil in tokamaks, which can be used to create an umbilic edge. Unlike tokamak divertor coils, an umbilic coil is non-axisymmetric and can have a fractional helicity. In the following sections, we discuss two cases, categorised by the direction of the current in the umbilic coils with respect to the toroidal plasma current.

5.1. *A US111 with an umbilic coil current opposite to the plasma current (counter-I case)*

Umbilic coils differ from helical coils by creating a ridge that may serve as a potential location for an X-point. These coils are similar to tokamak divertor coils, albeit without axisymmetry. To demonstrate their applicability, we modify an $n = 1$ MHD-stable equilibrium from the HBT-EP experiment ($R_0 = 0.94\,\mathrm{m}$, $A = 7.5$) by adding an $m = 1, n = 1$ ridge to it, breaking its axisymmetry, and then design an umbilic coil for it without modifying the toroidal or vertical field coils. The specific details of the equilibrium used are given in Appendix D.

https://doi.org/10.1017/S0022377825100925 Published online by Cambridge University Press

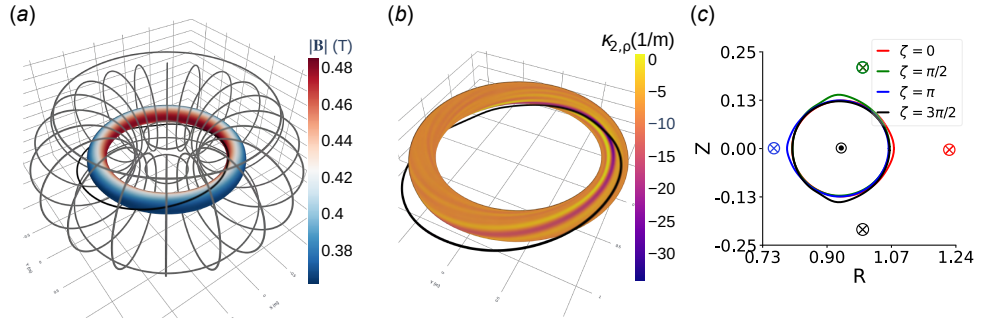


FIGURE 9. Output from the stage-two optimisation and free-boundary solution. (*a*) The full coilset of the modified HBT-EP experiment with the magnetic field strength on the boundary. (*b*) The umbilic coil and the curvature $\kappa_{2,\rho}$ on the boundary. (*c*) The cross-section of the LCFS from the free-boundary DESC solver along with the umbilic coil at four different toroidal angles and the directions in which they carry current relative to the plasma.

To accomplish this, we start by using 20 toroidal field (TF)[4] coils and 4 vertical field poloidal (PF) coils as done in the original HBT-EP experiment (Gates 1993; Sankar *et al*. 1993) and add an $m = 1, n = 1$ umbilic coil outside the plasma boundary at a distance of 1.5 times the minor radius. Similar to § 4.2, we perform second-stage optimisation to find the right shape of the umbilic coil. To ensure that the coils from stage-two optimisation give expected results, we calculate the free-boundary solution and present it in figure 9.

The effect of the umbilic coil can be decomposed as a rigid shift of the whole equilibrium and a small deformation of the boundary. Since the rigid shift (Sengupta *et al*. 2024) and the size of the ridge are small compared with the minor radius, it does not compromise omnigenity (as seen later in figure 12*a*). For a perfectly sharp umbilic edge, the magnetic fieldline must coincide with the edge, but even when the ridge lacks perfect sharpness, the magnetic fieldline stays close to it due to its high curvature. The curvature along the umbilic ridge $\kappa_{2,\rho} \in [-15, -32]\ \mathrm{m}^{-1}$ whereas the average curvature $\mathrm{avg}(\kappa_{2,\rho}) = -7\ \mathrm{m}^{-1}$.

The TF coils each carry a current of 105 kA, the PF coils carry 29.8 kA and the umbilic coil has a current $I_{\mathrm{umbilic}} = 2.1$ kA. But due to the proximity of the umbilic coil compared with TF and PF coils, it can still affect the plasma near the boundary significantly. As we can see, the umbilic coil is able to create an $m = 1, n = 1$ high-curvature ridge. The ridge formation is entirely attributed to the umbilic coil, as all other coils are axisymmetric. The current flowing through the coil is in the opposite direction to the plasma current. However, due to the three-dimensional nature of this set-up, the plasma still bulges towards the coil.

To understand the structure of the magnetic field near the edge, we use the FIELDLINES (Lazerson *et al*. 2016) code from the STELLOPT (Lazerson *et al*. 2020) optimiser. We observe the presence of a stochastic layer outside the plasma

[4]In the HBT-EP experiment, the TF coils have been wound multiple times and packed into a square aluminium frame with finite thickness of over 0.15 m, whereas we assume filamentary TF coils in our calculations, ignoring the effect of finite thickness. The radius of the TF filamentary coils is 0.65 m, which is the average of the outer and inner dimensions of a real HBT-EP TF coil.

https://doi.org/10.1017/S0022377825100925 Published online by Cambridge University Press

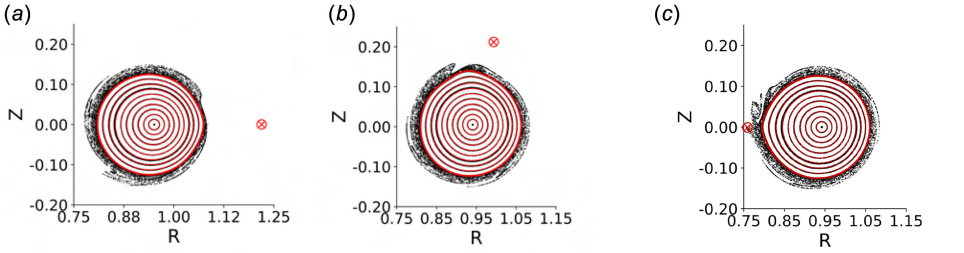


FIGURE 10. Poincaré section at different toroidal angles compared with the free-boundary DESC solution (red). The red cross marks the position of the umbilic coil.

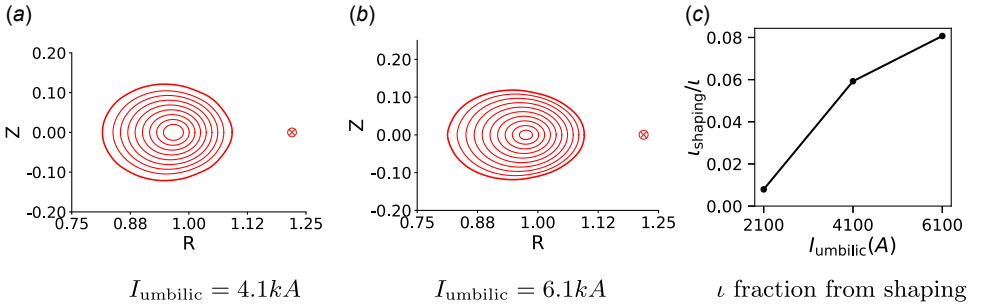


FIGURE 11. Free-boundary solution at $\zeta = 0$ for (a,b) different values of the umbilic coil current I_{umbilic} and (c) the fraction of the rotational transform generated by the umbilic coil current.[5]

boundary, similar to heliotron/torsatron configurations. Since the plasma and the umbilic coil are carrying currents in opposite directions, the poloidal field cannot vanish between them, which is why we do not observe an X-point. Therefore, the umbilic coil can create and vary the curvature of the ridge without creating and varying the current in the umbilic coil without creating an X-point. This also affects the fraction of the rotational transform generated by the umbilic coils. These details are provided in figure 11.

By controlling the current in the umbilic coil, we can shape the plasma without moving the limiter. Also, the ability to vary ι would help to stabilise various MHD instabilities and the $m = 1, n = 1$ tearing mode as done in the RFX-mod experiment (Piron *et al*. 2016) or experiments like the Madison Spherical Torus, where the edge rotational transform close to one can be achieved (Hurst *et al*. 2022). The J-TEXT (Liang *et al*. 2019) tokamak has recently proposed a similar idea (Li *et al*. 2024) of using umbilic coils to stabilise the edge of a tokamak.

[5] Since all the coils except the umbilic coil are axisymmetric and circular, any departure from a circular plasma shape is a consequence of the umbilic coil. Furthermore, for a perfectly circular equilibrium ($\iota_{\text{shaping}} = 0$), all the rotational transform on the boundary must come from non-circular shaping. Therefore, we conclude that any finite ι_{shaping} is a consequence of the umbilic coil.

https://doi.org/10.1017/S0022377825100925 Published online by Cambridge University Press

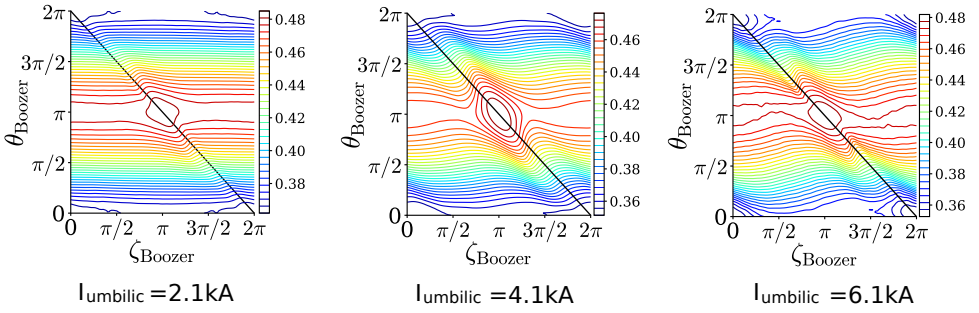


FIGURE 12. Magnetic field strength in Boozer coordinates on the boundary ($\rho = 1.0$) for different values of the umbilic coil current and the fieldline $\alpha = 0$. As the curvature of the ridge increases, so does the localised distortion of the $|\boldsymbol{B}|$ contours close to the fieldline.

Upon breaking the axisymmetry of the magnetic field, one must be careful as increasing the current I_{umbilic} for the same coil shape would degrade omnigenity and hence neoclassical transport. To understand the effect of umbilic coils on omnigenity, we present Boozer plots for the three values of current in figure 12.

We can see that the omnigenity degrades as we increase I_{umbilic} but the neoclassical losses are still acceptable: $\epsilon_{\text{eff}}^{3/2} < 0.001$. To control instabilities and transport due to MHD modes, we may need to generate a higher ι_{shaping} while ensuring omnigenity. One way to explore multiple such coil and equilibrium configurations would be to perform a single-stage optimisation, which simultaneously changes the plasma boundary and coil shapes. This is beyond the scope of the current study.

5.2. *A US111 with an umbilic coil current in the same direction as the plasma current (co-I case)*

In a manner similar to that of the previous section, we create a set-up where we modify the HBT-EP experiment and add an $m = n = 1$ ridge to the boundary. However, we do this while ensuring that the toroidal component of the current in the umbilic coil I_{umbilic} is in the same direction as the plasma current $I = -35.2\,\text{kA}$. To find the umbilic coil shape, coil current and the corresponding plasma shape that minimises the normal field error $\boldsymbol{B}\cdot\hat{\boldsymbol{n}}$ and the total pressure across the plasma boundary, we combine the objectives defined in (3.8) and (4.1) in a process known as single-stage optimisation. We solve

$$\min(\mathcal{F}_{\text{stage-one}} + \mathcal{F}_{\text{stage-two}}), \quad \text{s.t. } \nabla\left(\mu_0 p + \frac{B^2}{2}\right) - \boldsymbol{B}\cdot\nabla\boldsymbol{B} = 0, \quad \psi_{\text{b}} = \psi_{\text{b0}},$$
$$p = p_0(\psi), \quad I = I_0(\psi), \tag{5.1}$$

where ψ_{b0} is the user-provided enclosed toroidal flux by the boundary and $p_0(\psi)$ and $I_0(\psi)$ are the pressure and toroidal current profiles. The optimisable set of parameters is $\hat{\boldsymbol{p}} \in \{R_{\text{b},mn}, Z_{\text{b},mn}, a_k, X_n, Y_n, Z_n, I_{\text{umbilic}}\}$, where X_n, Y_n, Z_n are the Fourier coefficients that determine the shape of the umbilic coil. The coilset and free-boundary solution are presented in figure 13. After optimisation, the normalised $\max(\boldsymbol{B}\cdot\hat{\boldsymbol{n}}) < 0.5\,\%$ and the total pressure difference is $< 0.1\,\%$. Because we allow the boundary shape to change, we have less control over the curvature along the umbilic edge, so there is a lot of variation. The curvature along the edge lies in

https://doi.org/10.1017/S0022377825100925 Published online by Cambridge University Press

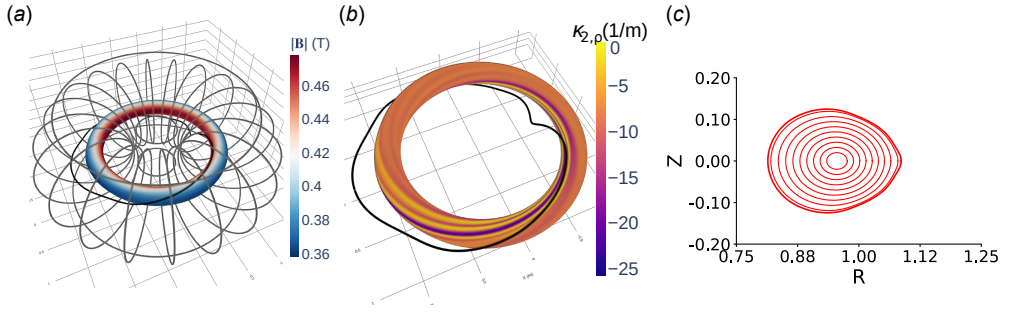


FIGURE 13. Coilset and free-boundary solution for the case where the current in the umbilic coil is in the same direction as the plasma current. The current carried by the umbilic coil is $I_{\text{umbilic}} = -3.1$ kA. All coils but the umbilic coil are axisymmetric.

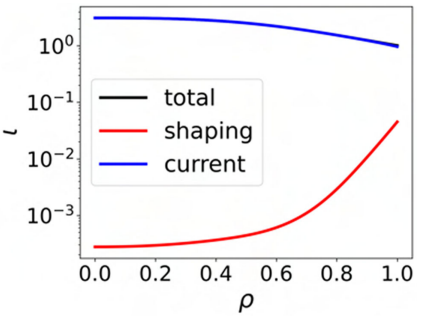


FIGURE 14. The total rotational transform profiles and the contribution from shaping and plasma current. The non-axisymmetric shaping is purely a consequence of the umbilic coil which contributes to approximately 5 % of the total rotational transform in the edge.

the range $\kappa_{2,\rho} \in [-25, -13]\,\text{m}^{-1}$ with the average curvature $\text{avg}(\kappa_{2,\rho}) = -8\,\text{m}^{-1}$. The neoclassical transport coefficient ${\epsilon_{\text{eff}}}^{3/2} < 0.001$ implies that the magnetic field is still omnigenous. Finally, we calculate the rotational transform profile of the optimised equilibrium to quantify the effect of the umbilic coil in figure 14.

The effect on the rotational transform profile due to the umbilic coil rises sharply as we move towards the edge. Therefore, it is possible that outside the plasma boundary the umbilic coil is able to affect the total magnetic field in such a way that it forms an X-point.

6. Conclusions and future work

This paper has explored the physical properties of USs. In § 3, we defined the umbilic parametrisation, approximated it using DESC and solved the ideal MHD equation inside the enclosed volume. In § 4, we developed vacuum and finite-β stellarators by simultaneously optimising the umbilic edge and the boundary shape using the DESC optimisation suite. We designed coils for the vacuum case, and traced field-lines in the edge of the US131 stellarator before and after adding a current source at the magnetic axis, to simulate a plasma fluctuation. Even though we do not see

https://doi.org/10.1017/S0022377825100925 Published online by Cambridge University Press

a perfectly sharp X-point (X-line) and well-defined fieldline structure like in a tokamak, the high-curvature edges could be a favourable location to place divertors and are robust to internal plasma fluctuations. We found that the optimised omnigenous umbilic designs do not possess a specific helicity regardless of the low-ripple transport, which could be a way to generate piecewise-omnigenous equilibria. As a final demonstration, in § 5, we propose an experiment to modify the Columbia HBT-EP experiment to an umbilic design by adding an $m = 1, n = 1$ helical coil. Using the DESC free-boundary solver, we demonstrate that the umbilic coils are able to cause the originally circular plasma to bulge and form a high-curvature ridge while maintaining omnigenity. By changing the direction of the current in the umbilic coil, we may be able to get an X-point, thereby converting the limited HBT-EP tokamak into a diverted stellarator.

The technique of simultaneously optimising a curve and a surface to control the magnitude and position of a sharp edge along a surface extends beyond USs. One can use this method to place and design ridges or regions of high curvature which may be useful for non-resonant divertors (Bader *et al*. 2018*b*), as well as these ridges creating preferential pathways for the fieldlines to exit the plasma (Bader *et al*. 2017). This can be achieved by multiplying a custom function $g(\theta)$ to the objective f_{umbilic} such that $g(\theta) = 1 \quad \forall \quad \theta \in \bigcup_i [\theta_{i0}, \theta_{i1}]$ and $g(\theta) = 0$ elsewhere along the umbilic curve, with θ_{i0}, θ_{i1} being optimisable parameters describing the lower and upper limits of regions where low curvature is penalised. This will create a set of continuous but disconnected ridges, just as we observe in non-resonant divertor concepts. This functionality is available in DESC. For an island divertor concept, the rotational transform on the boundary can be fine-tuned using umbilic coils, which could be used to reduce the sensitivity of X-points in the island divertor structure due to changes in the plasma boundary rotational transform. This idea is similar to that of divertor coils in tokamaks and would lead to a predictable strike point pattern.

Future work would involve using multiple umbilic coils to ensure plasma stability against low-n modes ($n = 0$ family) and increase the curvature of the ridge, similar to the work on tokamak separatrix by Webster (2009), understanding the strike point pattern and fieldline structure outside the plasma boundary or designing umbilic coils for a case similar to the finite-β US252 stellarator in § 4.4. It is also important to include MHD and kinetic stability metrics in the optimisation process, especially near the boundary as has been done for tokamaks by Bishop (1986) and Zheng *et al*. (2025) as stability properties can be sensitive close to the high-curvature plasma boundary.

Acknowledgements

One of the authors, R.G., enjoyed discussions with M. Zarnstroff, J. Levesque, R. Davies, J.L. Velasco and B. Jang, and gratefully acknowledges the help from S. Lazerson in installing the FIELDLINES code. We also express our gratitude to the referees for their constructive feedback that has improved the overall quality of this paper.

Editor Per Helander thanks the referees for their advice in evaluating this article.

Funding

This work is funded through the SciDAC programme by the US Department of Energy, Office of Fusion Energy Science and Office of Advanced Scientific

https://doi.org/10.1017/S0022377825100925 Published online by Cambridge University Press

Computing Research under contract no. DE-AC02-09CH11466, the Hidden Symmetries grant from the Simons Foundation/SFARI (560651) and Microsoft Azure Cloud Computing mini-grant awarded by the Center for Statistics and Machine Learning (CSML), Princeton University. This research also used the Della cluster at Princeton University.

Declaration of interests

M.L. is a consultant for Type One Energy.

Data availability statement

All the data, analysis and post-processing files are freely available in Gaur (2025). https://zenodo.org/records/15355215.

Appendix A. Calculating principal curvatures

The second principal curvature at a given point on a surface is defined as

$$\kappa_{2,\rho} = \min\left\{k : \det\begin{bmatrix} L_{\mathrm{sff},\rho} - kE & M_{\mathrm{sff},\rho} - kF \\ M_{\mathrm{sff},\rho} - kF & N_{\mathrm{sff},\rho} - kG \end{bmatrix} = 0\right\}, \tag{A1}$$

where

$$\begin{gathered} E = \frac{\partial \boldsymbol{x}}{\partial \zeta}\cdot\frac{\partial \boldsymbol{x}}{\partial \zeta}, \quad F = \frac{\partial \boldsymbol{x}}{\partial \zeta}\cdot\frac{\partial \boldsymbol{x}}{\partial \theta}, \quad G = \frac{\partial \boldsymbol{x}}{\partial \theta}\cdot\frac{\partial \boldsymbol{x}}{\partial \theta}, \\ L_{\mathrm{sff},\rho} = \frac{\partial^2 \boldsymbol{x}}{\partial \theta^2}\cdot\frac{\nabla\rho}{|\nabla\rho|}, \quad M_{\mathrm{sff},\rho} = \frac{\partial^2 \boldsymbol{x}}{\partial \theta\,\partial \zeta}\cdot\frac{\nabla\rho}{|\nabla\rho|}, \quad N_{\mathrm{sff},\rho} = \frac{\partial^2 \boldsymbol{x}}{\partial \zeta^2}\cdot\frac{\nabla\rho}{|\nabla\rho|} \end{gathered} \tag{A2}$$

are the metric coefficients corresponding to the first and second fundamental forms. The user-specified values $\kappa_{2,\mathrm{bound1}}$, $\kappa_{2,\mathrm{bound2}}$ define the limits on the curvature. Using this objective, we impose a maximum and minimum value on $\kappa_{2,\rho}$ that helps us impose a sharp curvature along the umbilic edge.

Appendix B. Form of various objectives

All the objectives with bounds are defined as

$$f_{\mathrm{obj}} = \mathrm{ReLU}(\max(\mathrm{obj} - \mathrm{obj}_{\mathrm{bound1}}, \mathrm{obj} - \mathrm{obj}_{\mathrm{bound2}})). \tag{B1}$$

In this paper, such a term is used to penalise the aspect ratio, the rotational transform and the second principal curvature $\kappa_{2,\rho}$ along the umbilic edge.

For the effective ripple objective explained in Unalmis *et al*. (2024) and the omnigenity objective implemented by Dudt *et al*. (2024), we do not use any bounds and try to reduce the magnitude of these quantities as much as possible. The exact definitions are given in the respective papers.

Appendix C. Piecewise omnigenity in umbilic configurations

In this appendix, we plot the normalised second adiabatic invariant $\mathcal{J}_\parallel = \int_{l_1}^{l_2} \mathrm{d}l\, v_\parallel/(2\pi R_0)$ as a function of the fieldline label α and the inverse pitch angle $1/\lambda \equiv E/\mu$. We calculate $\mathcal{J}_\parallel(\alpha, 1/\lambda)$ on a single flux surface on which the neoclassical transport coefficient $\epsilon_{\mathrm{eff}}^{3/2} \leqslant 0.001$. This is done for both the US131 vacuum and US252 finite-β equilibria presented in §§ 4.1 and 4.4, respectively. To highlight their

https://doi.org/10.1017/S0022377825100925 Published online by Cambridge University Press

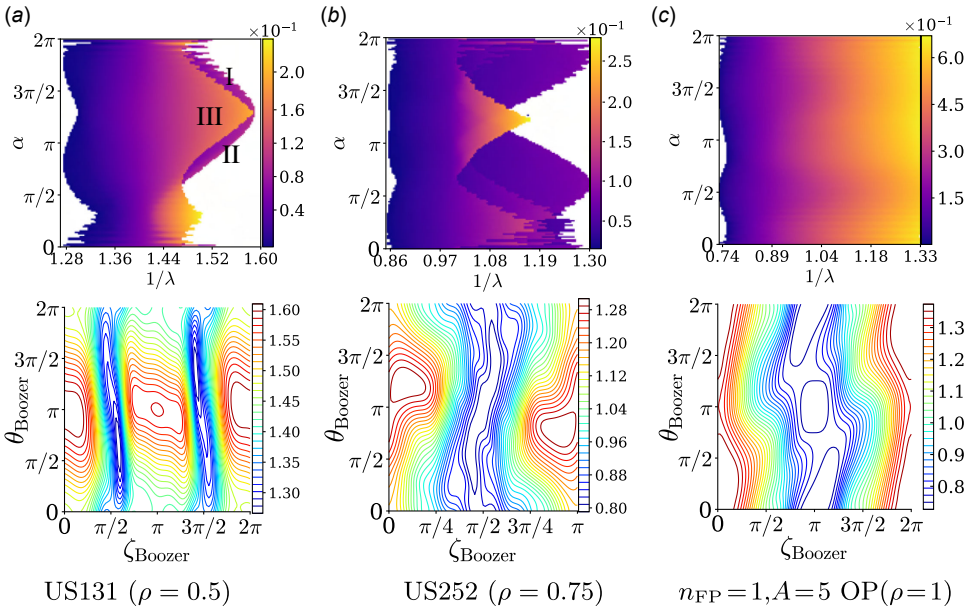


FIGURE 15. Plots of the second adiabatic invariant $\mathcal{J}_\parallel(\alpha, 1/\lambda)$ and $|B|$ contours for three equilibria with $\epsilon_{\mathrm{eff}}^{3/2} \approx 0.001$. For a close-to-perfect omnigenous equilibrium $\partial_\alpha \mathcal{J}_\parallel$ is close to zero for most values of $1/\lambda$, as seen in (c). However, for (a, b) there are multiple distinct values of $\mathcal{J}_\parallel$, discontinuous in α, a necessary condition for piecewise omnigenity.

difference, we also plot $\mathcal{J}_\parallel$ for a poloidally omnigenous (or QI) $n_{\mathrm{FP}} = 1$ vacuum equilibrium taken from the DESC omnigenity database (Gaur *et al*. 2023). Comparison is shown in figure 15.

For a perfect, non-piecewise-omnigenous stellarator $\partial_\alpha(\mathcal{J}_\parallel) = 0$, implying that $\mathcal{J}_\parallel$ is continuous and constant in α for almost every $1/\lambda$. Moreover, the distinct $\mathcal{J}_\parallel$ can be written as a linear sum such that the total $\mathcal{J}_\parallel$ for a value of the inverse pitch vanishes. However, for the umbilic configurations in this paper, especially the US131 vacuum stellarator, despite a low effective ripple (good omnigenity), we clearly see that $\mathcal{J}_\parallel$ is discontinuous in α for most of the inverse pitch values and $\tilde{\mathcal{J}}_{\parallel,\mathrm{I}} \approx \tilde{\mathcal{J}}_{\parallel,\mathrm{II}} \approx 0.5\tilde{\mathcal{J}}_{\parallel,\mathrm{III}}$. This supports our claim that the umbilic stellarator US131 is piecewise-omnigenous, especially at large values of the inverse pitch. US252 seems to be piecewise-omnigenous but the structure of $\tilde{\mathcal{J}}_\parallel$ is more complicated than that discussed by Velasco *et al*. (2024).

Appendix D. Details of the HBT-EP equilibrium

The original HBT-EP equilibrium is a shifted-circle equilibrium with a major radius $R_0 = 0.94$ m, minor radius $a_0 = 0.125$ m (aspect ratio $A = 7.52$) with $\beta = 0.5\,\%$ and a plasma current $I = -35.2$ kA. The pressure and rotational transform profiles are given in figure 16. The initial equilibrium was tested to be stable against $n = 1$ kink modes using the DCON-2D (Glasser 2016) stability code.

https://doi.org/10.1017/S0022377825100925 Published online by Cambridge University Press

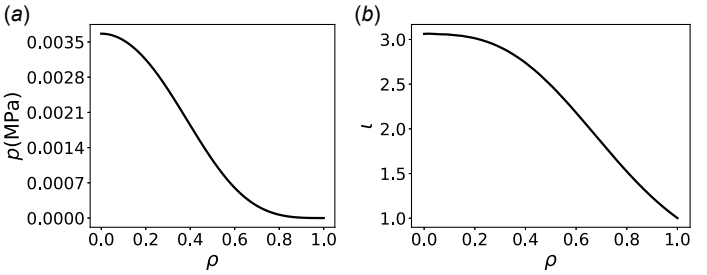


FIGURE 16. Profiles of the initial HBT-EP equilibrium.

REFERENCES

BADER, A., BOOZER, A. H., HEGNA, C. C., LAZERSON, S. A. & SCHMITT, J. C. 2017 Hsx as an example of a resilient non-resonant divertor. *Phys. Plasmas* **24**, 032506.

BADER, A., EFFENBERG, F. & HEGNA, C. C. 2018*a* Progress in divertor and edge transport research for stellarator plasmas. In *The Japan Society of Plasma Science and Nuclear Fusion Research*

BADER, A., HEGNA, C. C., CIANCIOSA, M. & HARTWELL, G. J. 2018*b* Minimum magnetic curvature for resilient divertors using compact toroidal hybrid geometry. *Plasma Phys. Control Fusion* **60**, 054003.

BISHOP, C. M. 1986 Stability of localized MHD modes in divertor tokamaks – a picture of the H-mode. *Nucl. Fusion* **26**, 1063.

BOOZER, A. H. 1983 Transport and isomorphic equilibria. *Phys. Fluids* **26** (2), 496–499.

BOOZER, A. H. & PUNJABI, A. 2018 Simulation of stellarator divertors. *Phys. Plasmas* **25** (9), 092505.

BURNETT, C. R., GROVE, D. J., STIX, T. & WAKEFIELD, K. 1958 The divertor, a device for reducing the impurity level in a stellarator. *Phys. Fluids* **1**, 438.

CARY, J. R. & SHASHARINA, S. G. 1997 Omnigenity and quasihelicity in helical plasma confinement systems. *Phys. Plasmas* **4**, 3323–3333.

CONLIN, R., DUDT, D. W., PANICI, D. & KOLEMEN, E. 2023 The DESC stellarator code suite. Part 2. Perturbation and continuation methods. *J. Plasma Phys.* **89**, 955890305.

DAVIES, R., SMIET, C. B., PUNJABI, A., BOOZER, A. & HENNEBERG, S. A. 2025 The topology of non-resonant stellarator divertors. arXiv preprint arXiv: 2501.18293.

D'HAESELEER, W. D., HITCHON, W. N. G., CALLEN, J. D. & SHOHET, J. L. 2012 *Flux Coordinates and Magnetic Field Structure: A Guide to a Fundamental Tool of Plasma Theory*. Springer Science & Business Media.

DUDT, D. W. & KOLEMEN, E. 2020 DESC: A Stellarator Equilibrium Solver. *Physics of Plasmas* **27**, 102513.

DUDT, D. W., CONLIN, R., PANICI, D. & KOLEMEN, E. 2023 The DESC stellarator code suite part 3: quasi-symmetry optimization. *J. Plasma Phys.* **89**, 955890201.

DUDT, D. W., GOODMAN, A. G., CONLIN, R., PANICI, D. & KOLEMEN, E. 2024 Magnetic fields with general omnigenity. *J. Plasma Phys.* **90**, 905900120.

FENG, Y., KOBAYASHI, M., LUNT, T. & REITER, D. 2011 Comparison between stellarator and tokamak divertor transport. *Plasma Phys. Control Fusion* **53**, 024009.

FENG, Y., SARDEI, F., GRIGULL, P., MCCORMICK, K., KISSLINGER, J., REITER, D. & IGITKHANOV, Y. 2002 Transport in island divertors: physics, 3d modelling andcomparison to first experiments on W7-AS. *Plasma Phys. Control Fusion* **44**, 611.

FORTUNATO, D. 2024 A high-order fast direct solver for surface PDEs. *SIAM J. Sci. Comput.* **46**, A2582–A2606.

https://doi.org/10.1017/S0022377825100925 Published online by Cambridge University Press

GARCIA, K. A., BADER, A., FRERICHS, H., HARTWELL, G. J., SCHMITT, J. C., ALLEN, N. & SCHMITZ, O. 2023 Exploration of non-resonant divertor features on the compact toroidal hybrid. *Nucl. Fusion* **63**, 126043.

GATES, D. A. 1993 Passive stabilization of MHD instabilities at high β_n in the HBT-EP tokamak. Tech. Rep., Plasma Physics Lab, Columbia University, New York.

GAUR, R. 2025 Omnigenous umbilic stellarators: data and analysis files. https://zenodo.org/records/15355215

GAUR, R., DUDT, D., CONLIN, R., PANICI, D., UNALMIS, K. & KOLEMEN, E. 2023 Exploring the database of omnigenous stellarators with DESC. In *APS Division of Plasma Physics Meeting Abstracts, APS Meeting Abstracts*, vol. 2023, p. AP01.033.

GLASSER, A. H. 2016 The direct criterion of Newcomb for the ideal MHD stability of an axisymmetric toroidal plasma. *Phys. Plasmas* **23**, 072505.

GOODMAN, A. G., MATA, K. C., PLUNK, G. G., SMITH, H. M., MACKENBACH, R. J. J., BEIDLER, C. D. & HELANDER, P. 2023 Constructing precisely quasi-isodynamic magnetic fields. *J. Plasma Phys.* **89**, 905890504.

GRIGULL, P. *et al.* 2001 First island divertor experiments on the W7-as stellarator. *Plasma Phys. Control Fusion* **43**, A175.

HALL, L. S. & MCNAMARA, B. 1975 Three-dimensional equilibrium of the anisotropic, finite-pressure guiding-center plasma: theory of the magnetic plasma. *Phys. Fluids* **18**, 552–565.

HARRIS, J. H., CANTRELL, J. L., HENDER, T. C., CARRERAS, B. A. & MORRIS, R. N. 1985 A flexible heliac configuration. *Nucl. Fusion* **25**, 623.

HIRSHMAN, S. P. & HOGAN, J. T. 1986 ORMEC: a three-dimensional MHD spectral inverse equilibrium code. *J. Comput. Phys.* **63**, 329–352.

HURST, N. C. *et al.* 2022 Self-organized magnetic equilibria in tokamak plasmas with very low edge safety factor. *Phys. Plasmas* **29**, 080704.

LANDREMAN, M. 2017 An improved current potential method for fast computation of stellarator coil shapes. *Nucl. Fusion* **57**, 046003.

LANDREMAN, M. & PAUL, E. 2022 Magnetic fields with precise quasisymmetry for plasma confinement. *Phys. Rev. Lett.* **128**, 035001.

LAZERSON, S., SCHMITT, J., ZHU, C., BRESLAU, J., & STELLOPT Developers, All 2020 STELLOPT. 10.11578/dc.20180627.6

LAZERSON, S. A. *et al.* 2016 First measurements of error fields on W7-x using flux surface mappinga. *Nucl. Fusion* **56**, 106005.

LI, Y. *et al.* 2024 The design of the external rotational transform coil on the J-TEXT tokamak. *Fusion Eng. Des.* **206**, 114591.

LIANG, Y. *et al.* 2019 Overview of the recent experimental research on the J-TEXT tokamak. *Nucl. Fusion* **59**, 112016.

NOETHER, E. 1918 Invariante variationsprobleme. Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen. *Math.-Phys. Klasse* **1918**, 235–257.

NÜHRENBERG, J. & ZILLE, R. 1988 Quasi-helically symmetric toroidal stellarators. *Phys. Lett. A* **129**, 113–117.

OHYABU, N. *et al.* 1992 Helical divertor in the large helical device. *J. Nucl. Mater.* **196**, 276–280.

OHYABU, N. *et al.* 1994 The large helical device (LHD) helical divertor. *Nucl. Fusion* **34**, 387.

PANICI, D., CONLIN, R., DUDT, D. W., UNALMIS, K. & KOLEMEN, E. 2023 The DESC stellarator code suite. Part 1. Quick and accurate equilibria computations. *J. Plasma Phys.* **89**, 955890303.

PANICI, D. *et al.* 2025 Surface current optimization and coil-cutting algorithms for stage-two stellarator optimization. arXiv preprint arXiv: 2508.09321.

PARRA DIAZ, F. *et al.* 2024 Flexible stellarator physics facility. Tech. Rep., Princeton Plasma Physics Laboratory (PPPL), Princeton.

PEDERSEN, T. S. *et al.* 2019 First divertor physics studies in Wendelstein 7-X. *Nucl. Fusion* **59**, 096014.

PIRON, C. *et al.* 2016 Interaction of external $n = 1$ magnetic fields with the sawtooth instability in low-q RFX-mod and DIII-D tokamaks. *Nucl. Fusion* **56**, 106012.

PUNJABI, A. & BOOZER, A. H. 2022 Magnetic turnstiles in nonresonant stellarator divertor. *Phys. Plasmas* **29**, 012502.

https://doi.org/10.1017/S0022377825100925 Published online by Cambridge University Press

RYUTOV, D. D. 2007 Geometrical properties of a "snowflake" divertor. *Phys. Plasmas* **14**, 064502.

SANKAR, M. K. V. *et al.* 1993 Initial high beta operation of the HBT-EP tokamak. *J. Fusion Energ.* **12**, 303–310.

SENGUPTA, W., NIKULSIN, N., GAUR, R. & BHATTACHARJEE, A. 2024 Asymptotic quasisymmetric high-beta three-dimensional magnetohydrodynamic equilibria near axisymmetry. *J. Plasma Phys.* **90**, 905900209.

SPITZER, L., Jr. 1958 The stellarator concept. *Phys. Fluids* **1**, 253–264.

UNALMIS, K. E., GAUR, R., CONLIN, R., PANICI, D. & KOLEMEN, E. 2024 Spectrally accurate reverse-mode differentiable bounce-averaging operator and its applications. arXiv preprint arXiv: 2412.01724.

VALANJU, P. M., KOTSCHENREUTHER, M., MAHAJAN, S. M. & CANIK, J. 2009 Super-X divertors and high power density fusion devices. *Phys. Plasmas* **16**, 056110.

VELASCO, J. L., CALVO, I., ESCOTO, F. J., SÁNCHEZ, E., THIENPONDT, H. & PARRA, F. I. 2024 Piecewise omnigenous stellarators. *Phys. Rev. Lett.* **133**, 185101.

WAGNER, F., BECKER, G., BEHRINGER, K., CAMPBELL, D., EBERHAGEN, A., ENGELHARDT, W., FUSSMANN, G., GEHRE, O., GERNHARDT, J. & GIERKE, G. 1982 Regime of improved confinement and high beta in neutral-beam-heated divertor discharges of the asdex tokamak. *Phys. Rev. Lett.* **49**, 1408.

WEBSTER, A. J. & GIMBLETT, C. G. 2009 Magnetohydrodynamic stability of a toroidal plasma's separatrix. *Phys. Rev. Lett.* **102**, 035003.

WEBSTER, A. J. 2009 Techniques for studying the separatrix of tokamak plasmas. *Phys. Plasmas* **16**, 012501.

ZEEMAN, E. C. 1976 The umbilic bracelet and the double-cusp catastrophe. In *Structural Stability, the Theory of Catastrophes, and Applications in the Sciences*, Lecture Notes in Mathematics, vol. 525, pp. 328–366. Springer.

ZHENG, L., KOTSCHENREUTHER, M. T., WAELBROECK, F. L. & AUSTIN, M. E. 2025 X-point effects on the ideal MHD modes in tokamaks in the description of dual-poloidal-region safety factor. *Phys. Plasmas* **32**, 012501.

https://doi.org/10.1017/S0022377825100925 Published online by Cambridge University Press