J. Plasma Phys. (2026), *vol*. 92, E120 © 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/S0022377826102256

Designing a buildable optimised stellarator to confine electron–positron plasmas

Pedro Francisco Gil[1], Jason Smoniewski[2,3], Paul Huslage[1], Rogerio Jorge[4], Timo Thun[2], Elisa Buglione-Ceresa[1], Tristan Schuler[5], Stefan Fingl[1], Grégoire-Hubert Ducas[6] and Eve Virginia Stenson[1]

[1]Max-Planck-Institut für Plasmaphysik, Garching 85748, Germany
[2]Max-Planck Institute for Plasma Physics, Greifswald 17491, Germany
[3]Plasma Science and Fusion Center, Massachusetts Institute of Technology, Cambridge, MA 02139, USA
[4]Department of Physics, University of Wisconsin-Madison, Madison, WI 53706, USA
[5]SchulerTEC, Germany
[6]Department of Mechanical Engineering, McGill University, Montreal, QC 13 H3A 0B9, Canada
Corresponding author: Pedro Francisco Gil, pedro.gil@ipp.mpg.de

(Received 9 April 2026 UTC; revision received 17 July 2026 UTC; accepted 27 July 2026 UTC)

In this paper, the design of the plasma equilibrium and superconducting coils for the electrons and positrons in an optimised stellarator electrons and positrons in an optimized stellarator (EPOS), experiment is presented. With newly developed stellarator optimisation tools, including single-stage and stochastic optimisation, as well as high-temperature superconductor strain, this work demonstrates that it is possible to achieve key metrics for the buildability and confinement properties of the device. In particular, satisfactory quality of quasisymmetry and stellarator robustness is designed, and engineering requirements are met for eight different candidates. A feasibility study is presented that optimises multiple candidates for different plasma major radii and coil currents, as well as the best EPOS candidate to date, named C4_R19.

Key words: fusion plasma, plasma devices, plasma confinement

1. Introduction

In magnetically confined fusion, deuterium and tritium are heated in order to produce an electron–ion plasma that is kept inside a magnetic field. This, however, has been revealed to be a highly complicated task, partly due to neoclassical and turbulent transport. One of the key drivers of a series of instabilities found in these plasmas is mass asymmetry, as electrons and ions react on different time scales. In electron–positron pair plasmas, in contrast, it is predicted that turbulent transport is suppressed, meaning that the behaviour is quiescent compared with its fusion counterpart (Tsytovich & Wharton 1978; Helander 2014; Helander & Connor 2016). There has been little experimental exploration of pair plasmas with electrons and positrons. Developing new devices capable of confining these entities could provide further insights into astrophysical plasmas, such as those found near pulsars and neutron stars (Arons 1979). Designing an experiment suitable for reproducing

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 19 Nov 2026 at 05:20:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0022377826102256

an electron–positron plasma on Earth requires the conception of a magnetic field configuration that can confine the two species stably for a sufficiently long time. The stellarator configuration is a suitable candidate, as it is an inherently steady-state machine, and it does not require driving currents in the plasma. External three-dimensionally shaped coils fully determine the shape of the magnetic field. Recent work (Landreman & Paul 2022) has shown that it is possible to optimise stellarator equilibria for precise quasisymmetry across all the magnetic surfaces. Quasisymmetry is a measure of how well collisionless trapped particles are confined. Generating enough positrons to meet the density requirements for a pair plasma is challenging, meaning that any loss of positrons is to be avoided; by being quasisymmetric, electrons and positrons in an optimized stellarator (EPOS), is able to confine the particles for longer which in turn increases the effects of cyclotron cooling. This is desired as cooling down the particles by radiating the perpendicular energy allows for longer confinement times. This goes against typical fusion-related plasma experiments that want to heat the plasma. Assuming an optimised stellarator, this would mean that limitations on confinement time originate from collisionless transport driven by real-world device imperfections and/or from collisional or turbulent transport.

For the EPOS stellarator, an on-axis field of 2 T is targeted to achieve cooling of the plasma through cyclotron radiation within less than a second (Kennedy & Helander 2021). This makes rare earth barium copper oxide (ReBCO) high-temperature superconductors (HTS) a promising candidate to realise the EPOS magnet system. The ReBCO superconductors have a high in-field performance and a critical temperature which can be comfortably achieved with compact cryocoolers without using cryogenic liquids. Non-insulated (NI) coils made from ReBCO superconductors show high thermal and electrical stability (Hahn *et al*. 2011). Non-planar NI coils have been demonstrated at low fields (Huslage *et al*. 2024*a*; Kim *et al*. 2025). We aim to use this approach for the EPOS coils in same spirit as described in Baillod *et al*. (2025, 2026).

Another aspect that has seen significant progress is the specific field of stellarator coil optimisation. Many codes have been proposed to optimise and find three-dimensionally shaped coils, such as COILOPT (Strickler & Hirshman 2002), FOCUS (Zhu *et al*. 2017; Hudson *et al*. 2018) or REGCOIL (Landreman 2017). Initially, these codes employed a two-stage approach: the first stage optimised the equilibrium to meet plasma specifications, and the second stage shaped the coils to closely reproduce the magnetic field generated in the first step. This approach can lead to incompatibility problems where the equilibria are far too complex to be reproduced by any buildable coils. To circumvent these issues, recent work by Henneberg *et al*. (2021), Jorge *et al*. (2023) and Giuliani *et al*. (2022, 2023) has enabled us to combine the two-step approach into a single-stage stellarator optimisation.

The paper is structured as follows: § 2 presents both the physics and engineering requirements for EPOS, outlining the targets that guide the optimisation. Then, in § 3, the techniques used for the coils and plasma optimisation in order to achieve the various configurations within the physics and engineering requirements are presented. Section 4 presents the results obtained from eight candidate configurations and provides information on the decision-making process leading to a final configuration, named here C4_R19, which possesses the necessary properties to proceed with an engineering design. Finally, in § 5, the results are summarised, and ideas for future improvement and analysis of the EPOS device are presented as the project advances towards its assembly and manufacturing phase.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 19 Nov 2026 at 05:20:14, subject to the Cambridge Core terms of use , available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0022377826102256

2. EPOS requirements

In this section, we explain how EPOS differs from other fusion-targeted stellarators and how this affects the optimisation of physics and engineering objectives. Size and magnetic field strength are two main drivers for the unique design of EPOS.

2.1. *Physics requirements*

As mentioned previously, the goal of EPOS is to confine electrons and positrons in a stellarator and to observe collective effects from the plasma; EPOS is hence different from its fusion counterparts. It is expected that EPOS will confine $10^{10}-10^{11}$ positrons and electrons, which means that it is going to be a low β plasma ($\beta < 10^{-11}\,\%$). In order to achieve conditions suitable for observable collective effects, a sufficiently high a/λ_D ratio is required, where a is the device's minor radius and λ_D is the Debye length; such a ratio measures how large the device is in comparison with the characteristic shielding length scale, which scales as

$$\frac{a}{\lambda_D} = a\sqrt{\frac{2n_p e^2}{\epsilon_0 k_B T_p}} \sim \sqrt{\frac{N_p}{T_p R}}. \tag{2.1}$$

Here, R corresponds to the machine's major radius, ϵ_0 is the vacuum permittivity and k_B is the Boltzmann constant, N_p, n_p, T_p and e refer to the number, density, temperature and charge of the positrons, respectively. The EPOS device will rely on cyclotron radiation cooling of the particles to a range of approximately 0.1–1 eV; this means that the magnetic field on axis should reach $\approx$2 T. The EPOS device therefore requires minimising the major radius while maximising the magnetic field amplitude in order to reach low enough temperatures. For cyclotron radiation cooling to be efficient, confinement times of the order of 1 s are necessary (Pedersen *et al*. 2012). In addition to the two design choices mentioned previously, namely machine size and magnetic field amplitude, which directly affect the optimisation scheme, the existence of the so-called 'weave-lane' (WL) coils is also a consideration. The latter will be explained in the next section.

Firstly, regarding dimensions, early estimations of device size show that a desirable upper boundary on the plasma volume is 10 L (Pedersen *et al*. 2012). This translates to a first approximation of a major radius of 20 cm and a minor radius of 5 cm, much like the design of the MUSE stellarator (Qian *et al*. 2023). The use of these values in the optimisation will be explained in § 3. To fulfil the size and field requirements, high-temperature superconductors (HTS) are suitable candidates for the coils. However, this brings geometrical consequences, meaning that the coils are not allowed to have prohibitive binormal curvature or torsion in order not to compromise the brittle nature of the HTS. Finally, considering the choice of a quasisymmetric field to confine collisionless particles for a sufficiently long time to cool them down, the quasi-axisymmetric (QA) configuration appears to be the most suitable for fulfilling EPOS requirements, namely high field strengths at small dimensions, thereby avoiding high coil complexity. This can be achieved because QA configurations are notably simpler than their quasi-helically symmetric and quasi-isodynamic relatives, as they strongly resemble a tokamak.

The choice of a QA field together with HTS coils limits the magnitude of the rotational transform ι that is required for EPOS. Unlike fusion stellarators, where the profile of the rotational transform can have a role controlling effects such as the Shafranov shift, ballooning modes and Mercier stability (Hegna & Nakajima 1998),

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 19 Nov 2026 at 05:20:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0022377826102256

EPOS parameters	Value/type	Unit
$\|B\|$ on axis	2	T
Particle energy	0.1–1	eV
a/λ_D	>10	–
Density	10^{12}–10^{13}	m^{-3}
Plasma volume	10	L

TABLE 1. Summary of the main physics parameters of EPOS.

for EPOS, a flat rotational transform profile is desired. By keeping its value constant radially, it is possible to avoid hitting any low-order resonant modes that could give rise to magnetic islands breaking up the flux surfaces and damaging the confinement quality. Moreover, early optimisation attempts revealed that the upper boundary of iota is determined by the trade-off between the HTS strain targeted to remain below the threshold and a low quasisymmetric error. Table 1 summarises the primary physics targets for EPOS.

2.2. *Engineering requirements*

In this section, the main engineering challenges are outlined when assessing the feasibility of a machine of the same dimensions as EPOS. As mentioned in the previous section, HTS will be the material of choice for the coils, and keeping its mechanical integrity via the minimisation of the binormal curvature and torsion is of high priority. For that matter, it is essential to keep the strain of the winding pack for both torsion and binormal curvature below a threshold of 0.2 % (Huslage *et al*. 2024*b*). The HTS coils will consist of double pancakes, meaning two adjacent HTS tape stacks, which will be wound around a three-dimensional metal frame. At this stage, it is planned to mill the trenches in which the tapes will lie using a 5-axis CNC machine. Each coil will be demountable and have its own unit, as shown in figure 1. Each double pancake possesses a 1 mm wall in the middle. A single pancake is 3 mm wide and approximately 17 mm tall, assuming a 0.1 mm spacing between HTS layers for solder paste and a tape height of 0.1 mm. The WL coils will consist of either 2, 3 or 4 double pancakes. The trenches will have some additional slack in width to accommodate winding the coil; early estimates indicate that 0.5 mm of slack should be sufficient for comfortable winding.

The milling and winding processes introduce several inaccuracies that accumulate, resulting in an overall imprecision of the current path geometry, commonly referred to as coil perturbations. These deviations are a notable source of problems in modern stellarators due to manufacturing tolerances, such as W7-X (Andreeva *et al*. 2004) and NCSX (Strykowsky *et al*. 2009), as they can potentially lead to poor confinement, the formation of magnetic islands and possibly the destruction of flux surfaces. As a reference, W7-X was built with a magnetic field amplitude error tolerance of 10^{-4} at 2 mm coil deviation amplitude (δ_{W7-X}). Assuming the scaling is linear, this would lead to the following tolerances for EPOS:

$$\delta_{EPOS} = \frac{\delta_{W7-X} a_{EPOS}}{a_{W7-X}} = \frac{1e-3\times 0.1}{1} \approx 10^{-4}\text{m}. \tag{2.2}$$

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 19 Nov 2026 at 05:20:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0022377826102256

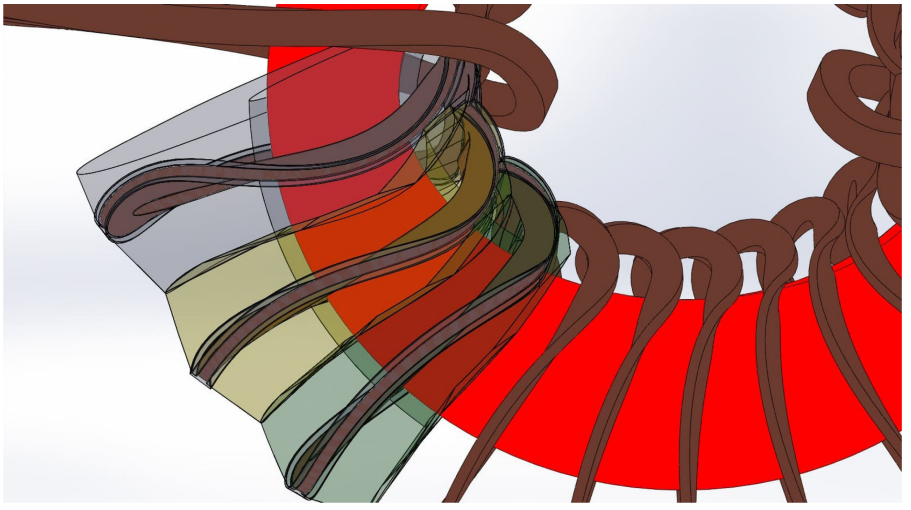


FIGURE 1. Top view of the engineering design concept for EPOS coils, each coil is made of a double HTS winding pack (here shown as a single brown stack) and is supported in an individual stellarator frame shown here in different colours for three coils. The last closed flux surface is shown in red.

Tolerances of 10^{-4} m, although not currently achievable with metal 3D printing technology, can be met via milling, which operates with sub-10^{-4} m precision. However, the extra slack mentioned previously, present in the width of the trenches, means that EPOS should still be able to present a precise enough QA field with a maximum deviation of 0.5 mm from trench uncertainty. Therefore, it is crucial to ensure that EPOS meets these robustness requirements, and one way to do so is through stochastic optimisation, which is described in § 3.3.

Another aspect that has to be taken into account when optimising EPOS are the $\boldsymbol{J} \times \boldsymbol{B}$ forces, where J is the current density passing through the coils and B the magnetic field. This can be of particular concern on the inboard side of the stellarator, where the coils find themselves closer to each other; an estimate for the Lorentz forces is shown in Appendix B at 1–2e5 N/m. These forces are induced by the other coils and the coil itself. If not considered, they can sometimes reach magnitudes high enough to deform metal structures, particularly in fusion devices. As a conservative measure EPOS will consist of a segmented shell, and it is on this shell that the trenches containing the HTS stacks take shape in several coils frames. As a result, EPOS will not have singular isolated modular coils, and the shapes of the coils will also ensure that the trenches remain accessible for winding. The coil frames are designed to allow for optimal alignment (minimising assembly errors) and conductive cooling of the HTS tapes.

Finally, the previously mentioned 'WL' coils are required for the injection of positrons via $\boldsymbol{E} \times \boldsymbol{B}$ drift, where E is the electric field applied by two electrodes near the confinement region, into the stellarator, similar to the injection technique into a magnetic dipole described in Stenson *et al*. (2018). These consist of two additional coils, also made of HTS, that are placed in toroidally opposed positions and possess a radius and current that are large enough to generate stray field lines not connected to the interior of the device, hence the name 'WL'. These field lines channel the positrons into two facing electrodes, inducing an electric field normal to the stray

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 19 Nov 2026 at 05:20:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0022377826102256

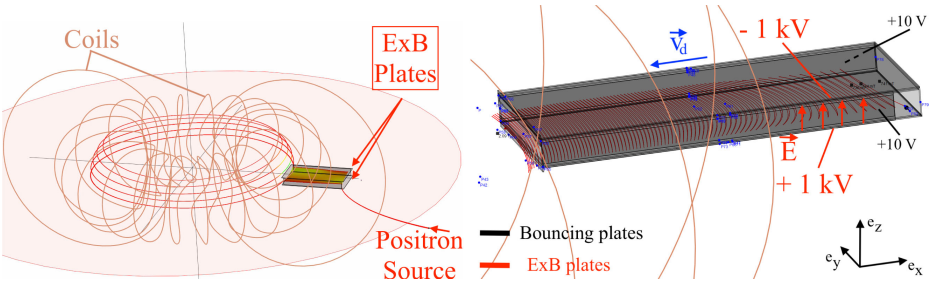


FIGURE 2. Left: injection concept for the EPOS stellarator, a single positron trajectory is plotted in red starting at a field strength of approximately 6 mT and then channelled by a stray field line originating from the WL coil. Right: the particles arrive at the E × B electrodes and experience a drift in the radial direction as well as multiple bounces (from additional bouncing plates) until it is finally injected into the three-dimensionally shaped field. The electrons are to be injected via an emitter.

magnetic field in order to force the positrons into the magnetic field containing the flux surfaces, as shown in figure 2. Additionally, two bouncing plates force the positrons into a back-and-forth motion as they drift radially inwards. Since they will be considerably larger than the standard coils, maintaining a magnetic field strength of 2 T on axis requires a larger current. At this stage of the design process, it is estimated that the standard coils will possess approximately 70 kA, and the WLs will carry even higher currents.

3. Optimisation method

Here, the goal is to present how we can integrate the physics and engineering requirements explained in the previous section into state-of-the-art stellarator optimisation and steer it towards an EPOS buildable configuration.

3.1. *Plasma objective function*

To address the challenge of designing EPOS, the task relies heavily on computational optimisation via the `SIMSOPT` framework (Landreman *et al*. 2021). It offers a broad set of tools to analyse, build and optimise stellarators. In our case, the challenge is twofold: find coils and an equilibrium that satisfy the aforementioned requirements. Regarding the physics part, `SIMSOPT` presents a Python interface with the Fortran Variational Moments Equilibrium Code (VMEC) (Hirshman & Whitson 1983), which minimises the magnetichydrodynamic (MHD) potential energy following a variational principle approach by assuming the existence of flux surfaces and so contains information on the plasma equilibrium, close to the STELLOPT optimisation suite (Lazerson *et al*. 2020)

$$W = \int \left(\frac{p}{\gamma - 1} + \frac{|B|^2}{2\mu_0} \right) \mathrm{d}V, \tag{3.1}$$

where W is the total MHD energy, B is the magnetic field, μ_0 is the vacuum magnetic permeability, p is the plasma pressure and γ is the adiabatic index. It can be shown that a state that minimises (3.1) solves the force balance equation

$$J \times B = \nabla p. \tag{3.2}$$

Here, J is the plasma current density. Although EPOS will not reach a regime where MHD is relevant, the VMEC code also allows for solving vacuum fields, and more

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 19 Nov 2026 at 05:20:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0022377826102256

importantly, finding configurations with good quasisymmetry. In this work, it is operated in fixed-boundary mode, meaning that the last closed flux surface (LCFS) is fixed, serves as a boundary condition, and the integral in equation (3.1) is calculated within its volume. The LCFS is described through a Fourier decomposition

$$R(\theta, \phi) = \sum_{m=0}^{M_{pol}} \sum_{n=-N_{tor}}^{N_{tor}} \left[a_{m,n} \cos(m\theta - n_{fp} n\phi) + b_{m,n} \sin(m\theta - n_{fp} n\phi)\right], \quad (3.3a)$$

$$Z(\theta, \phi) = \sum_{m=0}^{M_{pol}} \sum_{n=-N_{tor}}^{N_{tor}} \left[c_{m,n} \cos(m\theta - n_{fp} n\phi) + d_{m,n} \sin(m\theta - n_{fp} n\phi)\right], \quad (3.3b)$$

where M_{pol} and N_{pol} are the poloidal and toroidal modes, respectively, and n_{fp} is the number of field periods. The degrees of freedom that are optimised are the Fourier modes in the set of (3.3): $x_S = (a_{m,n}, b_{m,n}, c_{m,n}, d_{m,n})$. It should be noted that when the magnetic field presents stellarator symmetry (i.e. $B(R, -\theta, -\phi) = B(R, \theta, \phi)$), the sine terms and cosine terms can be skipped for R and Z, respectively. However, because of the WL coils, EPOS does not have stellarator symmetry, which increases the number of surface degrees of freedom to 625, with $M_{pol} = 12$ and $N_{tor} = 12$. These degrees of freedom are then adjusted during the optimisation to fulfil the following physics objective functions:

$$f_{QS} = \sum_{s_i} w_i \left\langle \left[\frac{1}{B^3} \left((N - \iota M) \boldsymbol{B} \times \nabla B \cdot \nabla \psi - (MG + NI) \boldsymbol{B} \cdot \nabla B \right) \right]^2 \right\rangle, \quad (3.4a)$$

$$f_\iota = (\iota - \iota_t)^2, \quad (3.4b)$$

$$f_A = (A - A_t)^2, \quad (3.4c)$$

$$f_\Delta = (\Delta - \Delta_t)^2, \quad (3.4d)$$

$$f_S = \left(\left\langle \frac{\mathrm{d}\iota}{\mathrm{d}s} \right\rangle - S_T \right)^2. \quad (3.4e)$$

Here, f_{QS} is a function penalising the departure from quasisymmetry, following the two-term formulation (Rodríguez *et al*. 2022). It is a weighted sum across all specified flux surfaces s_i (10 equally spaced in our case), with w_i its weights, B the magnetic field amplitude and (N, M) are parameters specifying the type of quasisymmetry aimed at, with (1, 0) corresponding to quasi-axisymmetry. Also, ψ is the toroidal flux function, and G is the poloidal and I the toroidal currents inside the surface. The parameter f_ι minimises the difference between the surface's rotational transform and the target set by the user; f_A and f_Δ perform a similar task but regarding the aspect ratio A $(=R/a)$ and the equilibrium radial width Δ of the bean cross-section. The latter is to prevent equilibria from violating the requirement of high a/λ_D. Finally, f_S is used to minimise the magnetic shear, giving a flat iota radial profile.

3.2. *Coils objective function*

For the coils, a method similar to the FOCUS code is presented, where coils are represented by geometric three-dimensional curves in space and their shapes are

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 19 Nov 2026 at 05:20:14, subject to the Cambridge Core terms of use , available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0022377826102256

varied during the derivative-based optimisation. Similar to the method presented for the plasma surface, coil number i is represented by a Fourier decomposition of the form

$$x^i(\theta) = \sum_{m=0}^{M_{coil}} x^i_{c,m} \cos(m\theta) + \sum_{n=0}^{M_{coil}} x^i_{s,n} \sin(n\theta), \tag{3.5}$$

and same for y^i and z^i. Note that θ here represents a coil length index and not the poloidal angle of the plasma. The order of decomposition is defined by the user. Equivalent to the plasma objective function, the coils' degrees of freedom are given by the Fourier modes of (3.5): $x_C = (x^i_{c,m}, x^i_{s,m}, z^i_{c,m}, z^i_{s,m})\ \forall i \in [1, N_{coils}]$ and $\forall m \in [0, M_{coil}]$. The EPOS device is planned to have 11 coils per field period including WL coils for 2 field periods and a total of 22 coils.

The degrees of freedom of the coils present in (3.5) are also complemented with an additional decomposition of the winding angle of the tape. This is done in order to reduce the strain applied to tape when winding (Huslage *et al*. 2024*b*). Using the same method defined in Singh *et al*. (2020), a centroid frame is constructed around the centre coil filament using a modified Frenet–Serret frame. The winding angle is defined as the angle between the normal vector $n_{centroid}$ in the centroid frame and the normal vector $\boldsymbol{n}$ of the actual HTS tape frame. This angle is then decomposed into a Fourier series in the following way:

$$\alpha(l) = \alpha_0 + \sum_{n=1}^{N_\alpha} \alpha_{c,n} \cos(2\pi n l) + \alpha_{s,n} \sin(2\pi n l), \tag{3.6}$$

where N_α is defined by the user and the coefficients $(\alpha_{c,n}, \alpha_{s,n})$ are additional degrees of freedom added to x_C. Note that α also coincides with the angle of the winding pack when considering the finite dimensions of the coils. Accounting for currents, $M_{coil} = 7$ and $N_\alpha = 3$, the total number of coil degrees of freedom amounts to 575. Since coil optimisation with three-dimensional filaments is an ill-posed problem (Imbert-Gérard *et al*. 2024), additional regularisation terms are added that allow for constraining the freedom of the optimiser when adapting the coil shapes. It is here that EPOS sets itself apart from conventional optimisation for stellarator symmetric devices, as a differentiation for the WL coils is necessary. Additional control over the shape, size and current of the WL coils is required, as they often prevent the optimiser from reaching good local minima. This is achieved by isolating specific functions in the optimisation, enabling the setting of specific weights for the WLs. The objective functions are therefore the following:

$$f_L^{WL} = \int_{\Gamma_0} \mathrm{d}l, \tag{3.7a}$$

$$f_L = \sum_{i=2}^{N_{coils}} \int_{\Gamma_i} \mathrm{d}l, \tag{3.7b}$$

$$f_{bin}^{WL} = \frac{1}{2} \int_{\Gamma_0} \max\left(\epsilon^0_{\mathrm{bin}} - \epsilon^{max}_{\mathrm{bin}}, 0\right)^2 \mathrm{d}l, \tag{3.7c}$$

$$f_{bin} = \sum_{i=2}^{N_{coils}} \frac{1}{2} \int_{\Gamma_i} \max\left(\epsilon^i_{\mathrm{bin}} - \epsilon^{max}_{\mathrm{bin}}, 0\right)^2 \mathrm{d}l, \tag{3.7d}$$

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 19 Nov 2026 at 05:20:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0022377826102256

$$f_{tor}^{WL} = \frac{1}{2}\int_{\Gamma_0} \max\left(\epsilon_{\text{tor}}^{0} - \epsilon_{\text{tor}}^{max}, 0\right)^2 \mathrm{d}l, \tag{3.7e}$$

$$f_{tor} = \sum_{i=2}^{N_{coils}} \frac{1}{2}\int_{\Gamma_i} \max\left(\epsilon_{\text{tor}}^{i} - \epsilon_{\text{tor}}^{max}, 0\right)^2 \mathrm{d}l, \tag{3.7f}$$

$$f_{msc}^{WL} = \frac{1}{L_0}\int_{\Gamma_0} \kappa^2 \,\mathrm{d}l, \tag{3.7g}$$

$$f_{msc} = \sum_{i=2}^{N_{coils}} \frac{1}{L_i}\int_{\Gamma_i} \kappa^2 \,\mathrm{d}l. \tag{3.7h}$$

Additionally, Γ is the coil path, L is the total length of the coil and N_{coils} is the number of coils per period. By convention, $i = 1$ corresponds to the WL coils. Here, the superscript f_L denotes the length of each coil, allowing for control over their size; l represents the coordinate of the coil length. Then, f_{bin} targets the binormal curvature of the HTS tape used in the coils, where ϵ_{bin}^{i} is the hard-way bending strain of coil number i, corresponding to when the HTS tape is bent around its strong axis, and ϵ_{tor}^{max} is equal to critical strain before the HTS tape breaks, in our case it is set to 0.2 % conservatively when compared with the literature (Paz-Soldan 2020) where a value of 0.4 % is given. The same applies to f_{tor} regarding torsional strain. Finally, the mean-squared curvature f_{msc} is also considered to prevent concave sections in the coils, which would significantly complicate the winding procedure of a single stack; here κ represents the local curvature.

3.3. *Stochastic objective function*

To relax the tolerances on the coil shapes, a single-stage optimisation approach was developed that utilises stochastic coil optimisation (Lobsien *et al*. 2018; Wechsung *et al*. 2022*a*,*b*; Gil *et al*. 2026*b*). Coil stochastic optimisation is specifically employed within the `SIMSOPT` framework. This method has been previously used for the design of the Columbia Neutral Torus, a stellarator comprising four coils: two interlocked and two poloidal field coils. There, stochasticity is considered by taking into account the tilts and translation errors of the coils. For EPOS, as described by Wechsung *et al*. 2022*a*,*b* the errors correspond to geometric deviations of the coil paths. The idea is to sample N random perturbed coils around the original coil and use them in the estimation of the field accuracy, measured through the quadratic flux metric, which will be explained in this section.

Considering the description of a coil as being a one-dimensional current-carrying filament of path $\gamma(l) = (X(l), Y(l), Z(l))$, with $l \in [0, 1]$ the coordinate along the coil, a new small perturbation $\epsilon_i(l)$ modelled as a Gaussian process (GP) is added to each component of γ

$$\tilde{\gamma}_i(l) = \gamma_i(l) + \epsilon_i(l). \tag{3.8}$$

The GP can be interpreted as a generalisation of random variables to functions, here, the function is one of the Cartesian coordinates $\gamma_i(l)$. It is described by a kernel $k(d)$, where $d = l_i - l_j$, which allows for the generation of a correlation function

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 19 Nov 2026 at 05:20:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0022377826102256

determining the regularity and smoothness of the perturbations along the coil. In this work, the radial basis function kernel is used, where:

$$k(d) = \sigma^2 \exp\left(\frac{-d^2}{2L_s^2}\right), \tag{3.9}$$

where $\sigma > 0$ measures the amplitude of the perturbations and $L_s > 0$ determines the characteristic length over which they occur. Moreover, in order to use this method in the optimisation routine, it is necessary to extract the derivatives of the GP and add them to the covariance matrix Σ. Then, after performing a Cholesky decomposition such that $\Sigma = DD^\top$ and by drawing a vector $X = (X_1, \ldots, X_n)$ from a multivariate normal distribution, where n is the number of quadrature points of the filaments, one vector of perturbation is generated by calculating $\epsilon_i = DX$.

Following this, N samples corresponding to N different perturbed stellarators are drawn. The perturbations are at first applied to the base coil set and multiplied via stellarator and field period symmetry, this is to reproduce repetitive manufacturing errors. Then, a second round of perturbations that break all symmetries is applied to all coils to take into account randomised production errors. Moreover, in the optimisation method employed here, the shaping of the coils is driven by the quadratic flux metric. It is a metric for the precision with which the coils reproduce the intended magnetic field

$$f_{SF} = \frac{1}{2}\int_S \frac{|\boldsymbol{B}\cdot\boldsymbol{n}|^2}{|B|^2}\mathrm{d}S. \tag{3.10}$$

Here S is the LCFS, $\boldsymbol{B}$ is the magnetic field generated by the coils, calculated through the Biot–Savart law, and $\boldsymbol{n}$ is the local normal to the surface, calculated from the plasma object. This term bridges the two objects that were previously presented, the coils and the plasma.

Finally, from the methods described above, it is therefore possible to compute N different quadratic fluxes from the N different stellarators. In a sample average approximation, the average squared flux is estimated about the single unperturbed VMEC surface object

$$\langle f_{SF}\rangle = \frac{1}{N}\sum_{i=1}^{N}\frac{1}{2}\int_S \frac{|\boldsymbol{B_i}\cdot\boldsymbol{n}|^2}{|B_i|^2}\mathrm{d}S. \tag{3.11}$$

The derivatives about the surface and coils' degrees of freedom used in the gradient calculations (Jorge *et al*. 2023) are therefore also the corresponding averages. It is worth noting that the stochasticity here is only applied in the estimation of the squared flux and not in the regularisation terms as they are not of relevance for the error calculations and could result in considerable discrepancies in the gradients amplitude (e.g. mean-squared curvature estimation) and prevent good convergence.

3.4. *Combined objective function*

Traditionally, stellarator design is performed in two steps: first, the plasma shape and properties are optimised, and then a second stage, in which suitable coils are found that reproduce the previous magnetic field. However, this method has an inherent issue: the equilibrium found in the first step may not be reproduced to desirable accuracy, given the engineering constraints. In this regard, recent work by

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 19 Nov 2026 at 05:20:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0022377826102256

Jorge *et al*. (2023) and tested in Baillod *et al*. (2025) combines the two steps into a single-stage approach. This enables both the plasma and the coils to be optimised simultaneously, and grants a higher compatibility between the coils and the equilibrium; this is the method of choice used to optimise EPOS. At this point, it is possible to write the minimisation problem as follows:

$$\begin{aligned} \min_{x_C, x_S} \quad & F(x_C, x_S) = F_S(x_S) + w_C F_C(x_C, x_S), \\ \text{s.t.} \quad & \psi = \psi_0, \\ & R = R_0. \end{aligned} \tag{3.12}$$

The w_C term is the weight that tunes the optimisation towards or away from coil targets to the detriment of surface objectives. In (3.12), the full objective function is a linear combination of the surface objective function F_S that is determined by the LCFS degrees of freedom x_S and the coils objective function F_C that is dependent on both the surface's and the coils' degrees of freedom since the quadratic flux is included in it. An additional constraint is applied to the parameter space, as the major radius is kept fixed throughout the optimisation. This also comes at a cost for the optimiser, which now has to deal with an increased number of degrees of freedom; this can impact the overall duration of the optimisation. For EPOS, when all the degrees of freedom are actively being optimised, it reaches a value of 1199.

3.5. *EPOS optimisation routine*

3.5.1. Initialisation parameters

As mentioned before, EPOS aims to confine electrons and positrons in a high magnetic field. Additionally, it is important to explore the optimisation landscape as thoroughly as possible. One choice that can affect this is the initial condition of the plasma; therefore, all the optimisation runs start with a simple plasma shape, also known as a 'cold start'. The LCFS shape is initialised with the coefficients $[a_{0,0}, a_{1,0}, a_{0,1}, a_{1,1}, d_{1,0}, d_{0,1}, d_{1,1}]$ being non-zero and to have a minor radius of $\approx$4 cm. These are chosen *ad hoc* so as to describe a simple low-order classical stellarator with very little shaping.

The coils are initialised before the optimisation as circular coils with a major radius 10 % larger than the plasma's, and with a minor radius at 70 % of the major radius of the plasma. However, additional care is required for the WL coils which start with a major radius 20 % bigger than the plasma's and a minor radius with the exact size of the major radius.

Moreover, the currents flowing in the coils are calculated to achieve an on-axis field of 2.0 T for a fixed number of coils. The number of coils for EPOS is determined through a study on the degree of quasisymmetric error versus the number of coils for an EPOS-like configuration (see Appendix A). The conclusion is that to achieve a satisfactory quasisymmetric configuration that is not affected by coil ripple, 10 coils plus one WL are required. This number, however, is also a strict upper boundary because at low enough major radius, the absolute coil-to-coil distance becomes too small, increasing the engineering challenge.

Additionally, the plasma primary radius degree of freedom is kept fixed during all steps of optimisation to ensure that it does not lead to a direct increase in the plasma volume nor bring the coils too close together.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 19 Nov 2026 at 05:20:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0022377826102256

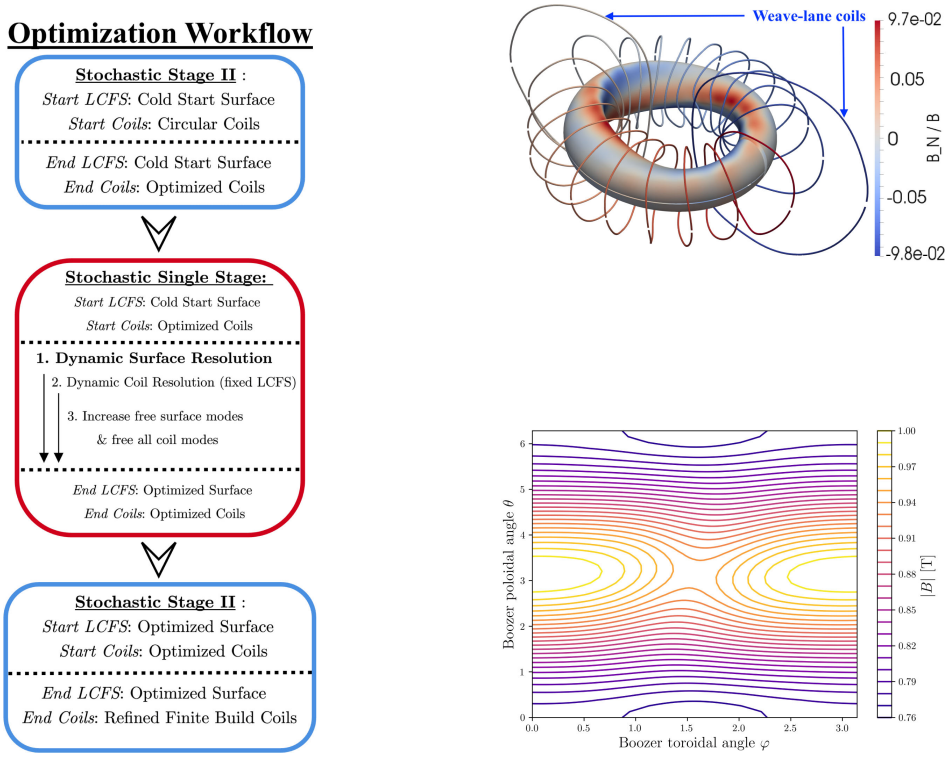


FIGURE 3. Left: diagram displaying the optimisation scheme for EPOS consisting of three main optimisation stages: first, a stage II optimisation is performed to find relatively accurate coils; followed by the stochastic single-stage method focused on improving the coil shapes and optimising the surface to balance physics and engineering constraints. Finally, stage II optimisation is required to refine the HTS strain and coil concavity. Top right: rendering of the EPOS configuration with LCFS and coils after the single-stage step, the heatmap corresponds to the field error $\langle B \cdot n \rangle / \langle B \rangle$, note that its maximum value is purposely still relatively high so as to not end up early on the design process in a sharp local minimum. Bottom right: Boozer plot showing the quality of quasisymmetry of the fixed-boundary 'cold-start surface' used for the EPOS optimisation.

3.5.2. Optimisation scheme

The goal of the optimisation is to produce an end configuration that exhibits good physical properties, such as a flat iota profile and good quasisymmetry, and with adequate engineering solutions, including feasible and robust coils that accurately reproduce the magnetic field. In figure 3, the workflow to achieve this is presented with all the starting conditions mentioned in the previous section factored in.

The procedure starts by employing a stochastic stage II optimisation similar to the work by Wechsung *et al*. (2022*b*) to get initial coils that match the 'cold-start' surface. This set of coils and surface can be visualised in figure 3. The objective function used for this preliminary optimisation is the following:

$$f_1 = w_{SF}\langle f_{SF}\rangle + w_L^{WL} f_L^{WL} + w_L f_L + w_{bin}^{WL} f_{bin}^{WL} + w_{bin} f_{bin} + w_{tor}^{WL} f_{tor}^{WL} + w_{tor} f_{tor}. \tag{3.13}$$

Then, using this configuration as a starting point, the stochastic single stage is executed with the same plasma optimisation technique used in stage I, sometimes

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 19 Nov 2026 at 05:20:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0022377826102256

referred to as dynamic surface resolution (Landreman & Paul 2022). Here, all the modes $[M_{pol}, N_{tor}]$ from (3.3) are fixed to start with, and then progressively one poloidal and one toroidal mode are unlocked after each iteration. Note that, within one mode iteration, one complete optimisation is performed. For example, in the first iteration, only the modes $m \leqslant 1$ and $|n| \leqslant 1$ are free. In this step, an analogous procedure is performed with the coil modes, which is referred to here as coil dynamic resolution. Once this phase is concluded, surface modes such that $m \leqslant 2$ and $|n| \leqslant 2$ are freed while keeping all the coil degrees of freedom also unfixed, and so on until all the coil degrees of freedom and all the surface degrees of freedom are unfixed. This method is adopted because it allows for a smooth transition of the coil shape. In contrast, the surface shape cannot change considerably, which prevents overloading of the optimiser with steep gradients in all directions. Freeing all the coil degrees of freedom for higher surface modes, which are unlocked, also enables sufficient flexibility in the coil shapes to match the increasing complexity of the surface as the dynamic resolution proceeds. In all the configurations presented in this paper, the full quadratic flux, as defined in (3.10), is targeted to be below 1e−4 after the single-stage step and to serve as a good starting point for the next stage. The plasma equilibria did not undergo a subsequent stage I after the single-stage step, as they already fulfilled the quasisymmetry and iota profile requirements for EPOS. In order to reach a high enough resolution in the surface description to account for coil ripple, the LCFS is optimised up until $M_{pol} = N_{tor} = 12$, as defined in (3.3), having in total 625 degrees of freedom, as mentioned previously. The coil filaments are generated with $M_{coil} = 7$ from (3.5), and $N_\alpha = 3$ in expression (3.6), which corresponds to 574 degrees of freedom or a total of 1199, including surface and coils. The objective function for the single-stage step is presented as follows:

$$\begin{aligned} f_2 = {} & w_{QS} f_{QS} + w_\iota f_\iota + w_A f_A + w_\Delta f_\Delta + w_S f_S \\ & + w_{SF} \langle f_{SF} \rangle + w_L^{WL} f_L^{WL} + w_L f_L \\ & + w_{bin}^{WL} f_{bin}^{WL} + w_{bin} f_{bin} \\ & + w_{tor}^{WL} f_{tor}^{WL} + w_{tor} f_{tor}. \end{aligned} \tag{3.14}$$

During the optimisation, the weight w_{QS} is progressively increased as adding enough degrees of freedom to resolve coil ripple ($N_{tor} \geqslant 10$) proves to deteriorate quasisymmetry when not enough weight is added. The target width in the optimisation is set to 4 cm for all the optimisations, as the volume should be minimised, and the target aspect ratio is set to 4.2. Moreover, the iota target is set to 0.2025, as higher values may require high HTS strain and coil complexity. It should be noted that these three functions have low weights compared with the other terms, as they are meant to vary and allow for mild constraint violation during the optimisation while still applying a barrier. The final step serves as a refinement step to ensure meeting the HTS strain engineering constraints, convex coils, finite build and size of the WL coils. This step is also stochastic, as all the previous ones, to ensure consistency in the methodology. The objective function used is

$$\begin{aligned} f_3 = {} & w_{SF} \langle f_{SF} \rangle + w_L^{WL} f_L^{WL} + w_L f_L \\ & + w_{bin}^{WL} f_{bin}^{WL} + w_{bin} f_{bin} \\ & + w_{tor}^{WL} f_{tor}^{WL} + w_{tor} f_{tor} \\ & + w_{msc}^{WL} f_{msc}^{WL} + w_{msc} f_{msc}. \end{aligned} \tag{3.15}$$

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 19 Nov 2026 at 05:20:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0022377826102256

Here, the coils are modified to have finite dimensions in the optimisation. This is achieved by simulating a winding pack with dimensions of 7 mm × 15 mm, consisting of 3 × 3 current-carrying filaments. Not only does it enable a higher-fidelity calculation of the actual magnetic field, but it also provides a visual assessment of the tape's complexity, including its twists and bends. Then, f_{msc} is also added so as to ensure coil convexity and prevent strong variations of the curvature. A considerable amount of fine tuning of the weights is required in this final step. As guidelines, the optimisation is set to simultaneously meet a total squared flux smaller than 4×10^{-6} m^2 using the local normalisation method, an HTS strain in all the coils lower than 0.2 %, negative curvature (for concavity reasons) no higher than -2 m^{-1} and WL coils no longer than 1.5 m and still displaying stray fields allowing for $\boldsymbol{E} \times \boldsymbol{B}$ injection. The value for the squared flux should be considered as an *ad hoc* condition from previously made robustness verifications. These results showed that configurations with a squared flux higher than 6×10^{-6} m^2 and coils perturbed with an amplitude higher than 0.5 mm do not yield sufficiently precise quadratic minimising (QFM) surfaces with values as high as 6×10^{-6} for the QFM residual.

3.5.3. Definition of buildability

The landscape of buildable stellarators meeting all the physics and engineering requirements mentioned previously is not trivial to find. In most optimisation problems that are not inherently convex, there can be an extensive number of local minima. To explore this parameter space, two configuration scans are performed: one for the major radius and one for the ratio between the current of the WLs and the one from the standard coils. Regarding the major radius, four values are explored: {16, 17, 18, 19} cm, which are fixed during optimisation. Lower values are not considered, as they bring the coils too close together to preserve a low squared flux, and higher values are ruled out since they naturally require more positrons. For the WL current ratio, two ratios are considered: {3, 4}, a ratio of 2 is not considered, as the optimiser consistently reduces the size of the WL coils in order to approximate as much as possible a stellarator symmetric device, hence preventing the presence of stray fields for WL coils with small dimensions. Note, however, that the values for the current ratio serve as initial conditions for the optimisations but do not act as an extra fixed degree of freedom. Giving still freedom for the optimiser to deviate from these initial values and enable access to perhaps additional configurations. Therefore, in this paper, eight different configurations are presented, all of which can be designed and built. All the stochastic optimisations are performed with an absolute perturbation amplitude σ of 1 mm, so as to overshoot the precision requirements of 0.5 mm in the trenches that are expected at the current engineering state. This is to account for additional tolerances that have to be taken into account due to assembly inaccuracies, thermal loads and electromagnetic forces. These perturbations are done with a characteristic length $L_s = 0.2$ m, given that most of the coil deviations will originate from the slack in the winding pack trenches, and higher frequency perturbations should not arise due to the tape's stiffness in the binormal and normal directions.

Moreover, the WL stray fields should correspond to a group of field lines extending from the plane of the WLs, crossing a boundary set by a cylinder of 60 cm in radius and approximately 80 cm high corresponding to the vacuum chamber walls, and reaching a region of 5 mT outside the chamber, taken as the standard field strength of guiding tubes necessary to inject the positrons; this is shown in figure 3. Another aspect considered in the optimisation and subsequently checked for all

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 19 Nov 2026 at 05:20:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0022377826102256

candidates is the presence of positive normal curvature. Normal curvature in the rotated centroid frame is defined as

$$\kappa = \frac{\tilde{\boldsymbol{N}} \cdot \tilde{\boldsymbol{T}}'}{\|\boldsymbol{\Gamma}'\|}. \tag{3.16}$$

Here, Γ describes the curve in three dimensions, $\tilde{\boldsymbol{T}}$ is the tangent vector to the curve and $\tilde{\boldsymbol{N}}$ is a modified normal curvature vector (see Huslage *et al*. (2024*b*)). This quantity is positive in concave sections of the coils. The tape will be similarly wound around the trenches, just like magnetic tape is wound on a spool. This means that any positive curvature sections would make the winding procedure considerably more complicated, as no tension can be sustained throughout the process.

The simulations are all performed on the VIPER cluster of the Max-Planck Computing and Data Facility, which comprises a total of 768 nodes featuring AMD EPYC Genoa 9554 CPUs with 128 cores and at least 512 GB of RAM per node. Stochastic optimisation depends on an MPI implementation of the codes running on five nodes with 128 cores each, where one core executes the code and computes the results for one sample. Together, the optimisation uses 1280 samples, and 350 GB of RAM is reserved for these purposes at an average of approximately 10 node-hours or 1280 core-hours for each stochastic single-stage run. For a single configuration the stochastic single-stage step is done approximately 20 times to allow for hyperparameter tuning.

4. Numerical results

4.1. *EPOS candidates*

The investigation of potential designs for the EPOS is explained here, with eight configurations of both coils and equilibrium. They are labelled as C(3,4)_R(16,17,18,19), where C refers to the standard coil to WL coil current ratio and R to the major radius of the configuration. The analysis presented in this paper is conducted to the extent necessary to allow for a first design freeze of the configuration. A subsequent engineering design phase is necessary to assess the feasibility of the stellarator fully. This will include, for example, inspections of the coil stresses, magnetic permeability of materials and their error fields with finite-element simulations. If any modifications are needed, other configurations may be considered or modified. The metrics that are shown here are calculated from the reconstructed VMEC equilibria generated by the coils. This can be achieved by optimising a QFM surface for the LCFS to fit the one actually produced by the coils. Then, by updating the boundary and rerunning VMEC, the information for the vacuum equilibrium can be obtained. The optimisation for all cases is performed using 16 poloidal modes and 22 toroidal modes, with a tolerance error of 10^{-12}, sufficient to resolve the coil ripple of the magnetic field lines, which can significantly damage the quasisymmetric error (QS error). The coils used in the reconstruction are generated with nine filaments per coil to simulate their finite dimensions, as described previously.

In figure 4, both the total QS error and the total squared flux as a function of the QS error are shown. The prior is the sum over 11 flux surfaces, and the total QS error is plotted for both the VMEC and the QFM values. Here, VMEC circle data points correspond to the target QS error value from the VMEC equilibrium object that the optimiser obtained in the single-stage step, which is not the one generated by the coils. In contrast, the diamond QFM data points represent the QS value from the reconstructed equilibria from the final coils. The first observation is that all the

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 19 Nov 2026 at 05:20:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0022377826102256

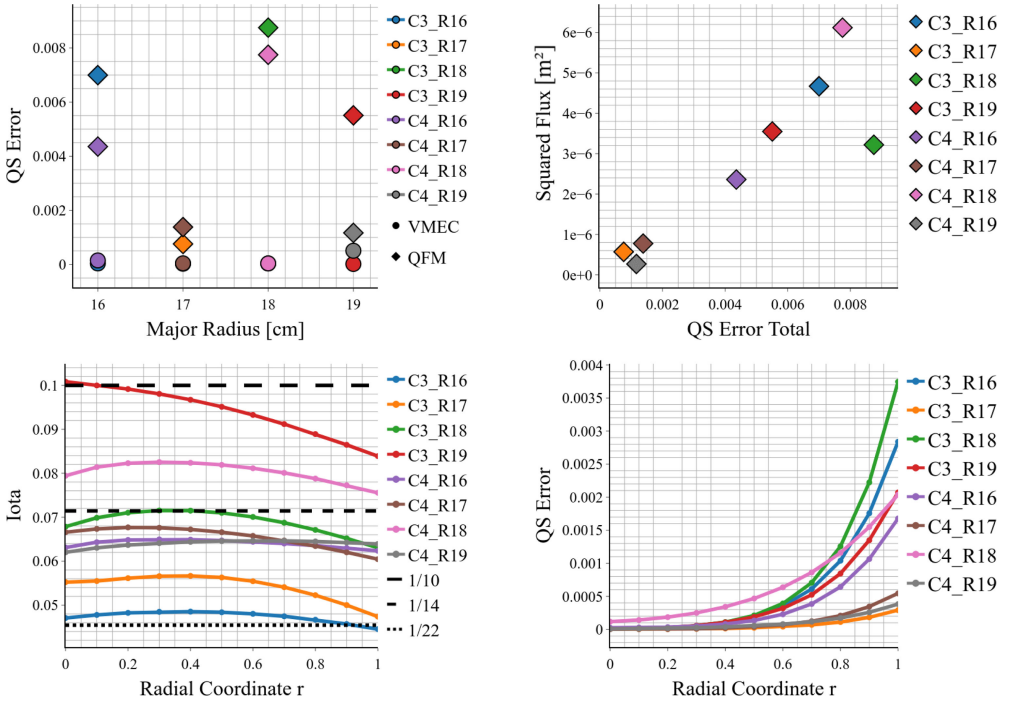


FIGURE 4. Top left: quasisymmetric error of the EPOS configurations plotted against the major radius for both the VMEC equilibrium and the QFM-generated one. Top right: squared Flux vs QS error for all configurations. Bottom left: iota profiles for all the EPOS candidates and the three most important rationales bounding the rotation transform values. Bottom right: quasisymmetric error radial profile for the EPOS configurations.

QFM values are higher than the VMEC ones. This is expected, as they primarily come from coil ripple, given that there is a finite number of coils with finite dimensions. A considerable variation in the QS quality of the QFM configurations is also visible, as the configurations are optimised to account for engineering constraints, which makes the choice of weights unpredictable and results in very different configurations, rather than just scaled-up or -down versions of each other. However, for 16, 18 and 19, all coil sets with a current ratio of 4 performed better than their counterparts with a current ratio of 3. The case C3_R17 can be considered an exception, as it is optimised with the purpose of disregarding concavity in the coils, ultimately allowing it to reach regions of high accuracy in the optimisation space more easily.

Moreover, as shown on the top right plot of figure 4, the better the total squared flux, the better the quasisymmetry of the considered configuration. Note that the correlation is not one-to-one, as local variations of the field generated by the coils on the target magnetic surface are only on average detrimental to the quasisymmetry, and a certain percentage of these can amount to a local improvement of quasisymmetry. However, the most accurate configurations are those with the lowest squared flux, which can result in an almost one order of magnitude better QS metric. Another aspect to consider when designing a stellarator is the iota profile. As mentioned previously, for EPOS, a flat iota profile value between 0 and 0.1 is targeted. As shown in the bottom left of figure 4, all the configurations besides C3_R19 possess an iota that varies at most 10 % radially. The highest ratios that bound the iota profiles,

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 19 Nov 2026 at 05:20:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0022377826102256

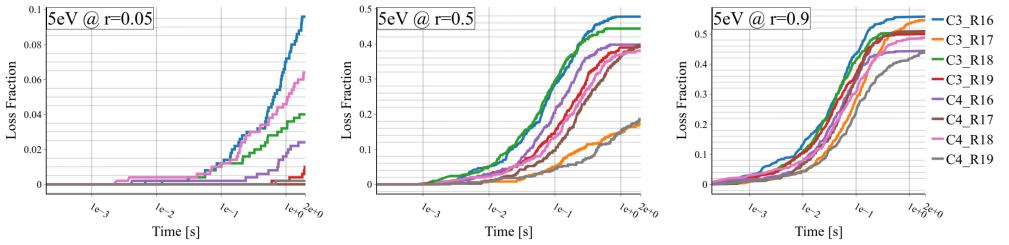


FIGURE 5. Loss fraction of positrons at 5 eV versus time of the EPOS configurations with starting conditions at $s = 0.05$ (left), $s = 0.5$ (middle) and $s = 0.9$ (right). Here, we show that, in the worst case scenario with no cyclotron cooling, at 5 eV, we are able to keep almost 100 % of the particles after 2 s when starting the particles at the magnetic axis. This ratio reduces to approximately 60 % when starting at the outer boundary.

except for C3_R19, are 1/14 and 1/22, meaning that no major island chains are expected to appear in the magnetic field.

The fact that the iota of all the configurations remained relatively low compared with other known devices, such as W7-X ($\iota \sim 1$) or MUSE ($\iota \sim 0.2$) (Pedersen *et al*. 2016; Qian *et al*. 2023), allows EPOS to find simpler coils with lower binormal and torsional strain that can potentially damage the coils during winding, for example. Then, although the possible values for the major radius are between 16 and 19 cm and the target aspect ratio is 4.2, the optimiser kept all of the aspect ratios below 4.2; on top of that, the bean cross-section width is also penalised with a target value of 4 cm, making it more compact than anticipated. This is related to the fact that the weights of the aspect ratio and width are purposely kept low in comparison with other terms, allowing more freedom for the gradient descent to find suitable minima. The variation in QS error can be further observed in the profiles on the bottom right of figure 4, where three configurations stand out as particularly accurate: C3_R17, C4_R17 and C4_R19 with values all lower than 10^{-3} throughout the volume, which is coherent with the results displayed on the top left of figure 4. Moreover, as expected, all configurations perform substantially better as they approach the magnetic axis. This suggests that, regardless of the configuration, core injection should be considered whenever possible to maximise the confinement time of positrons and electrons. This is supported by the data displayed in figure 5. The data were produced using the code SIMPLE (Albert *et al*. 2020*a*,*b*), which uses symplectic integration in systems where the Hamiltonian is specified in non-canonical coordinates for the estimation of alpha particle losses of guiding-centre orbits in a three-dimensional geometry, but adapted for the loss of positrons. The simulations are performed with 500 particles with 5 eV and a random distribution of $v_\perp/v$ and at different radial coordinates: 0.05 (at the magnetic axis), 0.5 and 0.9 (near the LCFS). This is done to inform on the confinement performance of all the configurations depending on the injection scheme.

The time duration of the particle simulation is 2 s, which is estimated to be long enough to radiate approximately 40 % of the perpendicular energy through cyclotron cooling (Anderegg, Hollmann & Driscoll 1998). This ensures that most of the particles possess mainly parallel energy, i.e. are passing particles and therefore can be indefinitely confined assuming flux surfaces and no other transport effects. It is visible that starting the simulations at the core yields, in the worst case scenario, a loss of only 10 % of particles after 2 s and, in the best case, 0 % losses. This confinement

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 19 Nov 2026 at 05:20:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0022377826102256

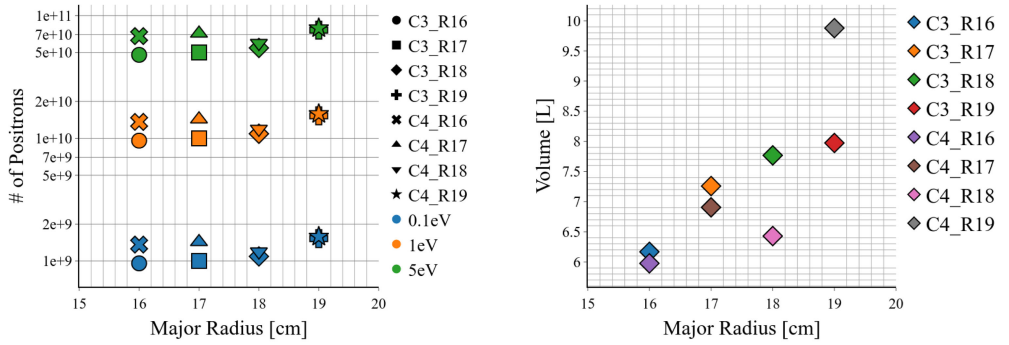


FIGURE 6. Left: minimal number of positrons required to reach $a/\lambda_D \geqslant 10$ for all the configurations as a function of the major radius and plotted for different temperatures (0.1, 1 and 5 eV). Right: volume taken by the vacuum field of all the configurations as a function of the major radius.

is, however, degraded as the particles are injected closer to the edge with losses reaching 50 % with two outliers (C3_R17 and C4_R19) at less than 20 %. Finally, for all configurations, results that start with the particles close to the LCFS simulating an edge-injection-type scenario reach between 55 % and 40 % losses. These results illustrate the quality of the confinement/quasisymmetry of the configurations; however, little can be concluded about the injection scheme for EPOS, as it will need to be fine tuned experimentally primarily. Moreover, the number of confined particles of the built machine will be degraded to some extent due to injection inefficiencies and manufacturing and assembly errors, which will further impact the quality of quasisymmetry.

The main idea of the scan is driven by the scaling of a/λ_D as a function of the major radius, as shown in (2.1), where a smaller major radius means a higher ratio and therefore it reduces the amount of required positrons. Figure 6 shows the minimal amount of necessary positrons to meet the requirement of $a/\lambda_D \geqslant 10$ for all the EPOS configurations as a function of the major radius and for three different positron temperatures, namely 0.1, 1 and 5 eV. It is visible that the temperature is the parameter that primarily impacts the number, and that the major radius of the configurations does not impact as much. Another quantity that can directly impact the a/λ_D ratio is the volume (shown to the right in figure 6). As expected the volume increases with the major radius of the configurations, meaning that the Debye length increases. Since there is such a high variability in parameters such as the minor radius and hence in the aspect ratio, as supported by figure 6, this can explain the minor enhancement of the necessary particle number and simultaneously means that the number of required positrons of a configuration is no longer a primary parameter in the choice of the stellarator.

Designing buildable coils has two main goals: simplification of the manufacturing process, by minimising the development of winding machines and winding pack technology, for example, and ensuring that the designed coils can be replicated within reasonable accuracy. This is because extremely twisted and complex coils are inherently more difficult to wind. Therefore, it is additionally laborious to accurately shape them, as well as to verify that they comply with tolerance requirements. In figure 7 are shown the maximum torsional and binormal strain values of all the configurations. All the values are below the 0.2 % threshold, which is already a conservative

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 19 Nov 2026 at 05:20:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0022377826102256

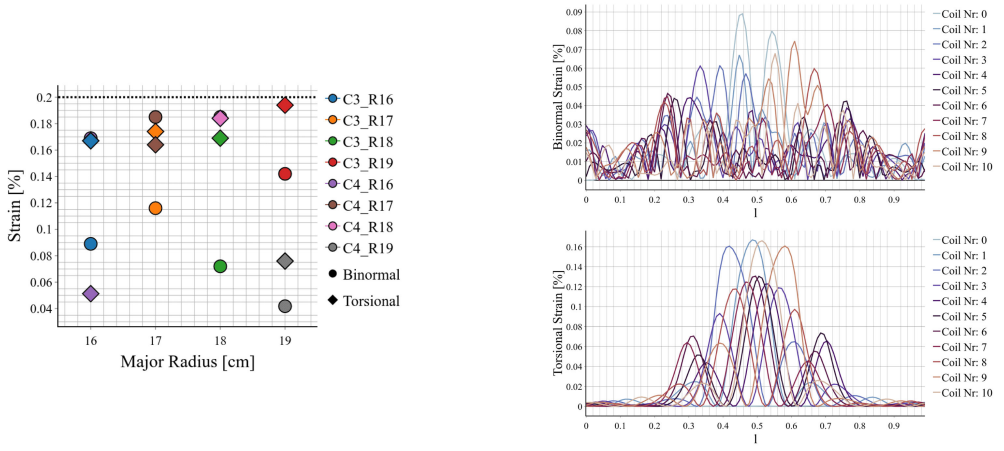


FIGURE 7. Left: maximum HTS torsional and binormal strain of all the configurations presented. Top right: HTS binormal strain of all the coils of the C3_R16 configuration. Bottom right: HTS torsional strain of all the coils of the C3_R16 configuration.

value for the maximal accepted strain value, while simultaneously showing acceptable squared fluxes that all score below 6×10^{-6} m^2, as shown in figure 4. An example of the strain profiles of what is considered a safe coil to build is shown on the right in figure 7. These profiles are characteristic of QA for EPOS, as the inboard side (at $l \approx 0.5$) requires the most shaping of the coils, which is evident in both torsional and binormal strain. The shaping can also be observed in the curvature κ profiles and the physical coils of figure 8. The areas of high coil complexity are, therefore, in the EPOS case, very localised.

Moreover, EPOS will be built using NI coils. As highlighted earlier, concave sections prevent the tension from being kept on the tape while winding. Tension is necessary as it establishes good thermal and electrical contacts between HTS layers in the winding packs. Thanks to the subsequent current and heat redistribution in case any section of the HTS becomes normally conducting, the coils become passively robust to quenches, which is why NI coils are of such interest in magnet research (Wang, Chan & Schwartz 2016; Hahn *et al*. 2016). Therefore, a requirement for EPOS is that all the coils are convex. In figure 8, the standard curvature profiles along each coil of configuration C4_R19 and of coil 1 of the set C3_R17 are shown. The curvature of all the coils of the former stellarator is negative, indicating a consistent convex structure of the coil, which is also visible in the three-dimensional model of coil 1 of the respective stellarator. As an example of what the optimiser systematically finds as a stable minimum, the curvature of coil 1 of C3_R17 is also plotted, displaying very strong concavity around the inboard side of the coil, reaching a value of 20 m^{-1}. To understand its concavity, the coil is also shown on the right. A visual difference between the two coils is that C3_R17 possesses the characteristic 'S' shape on the inboard side that gives rise to the rotational transform, which, together with a minimisation of the torsional and binormal strain, enhances the concavity of the coil. It is nevertheless possible to generate a rotational transform without twisting the coils to such complex levels and without breaking the HTS strain limits. This method will be explained in the following section.

Another essential parameter when considering building an EPOS-sized device is the manufacturing tolerances. This is the main reason why stochastic optimisation is

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 19 Nov 2026 at 05:20:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0022377826102256

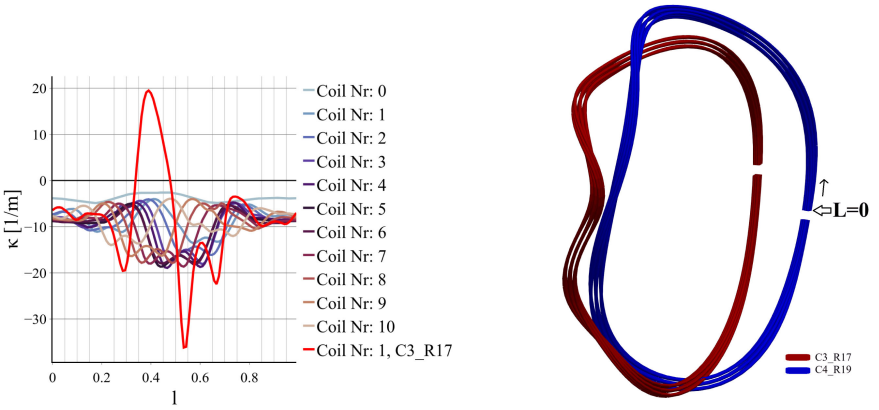


FIGURE 8. Left: normal curvature profiles for all the coils of configuration C4_R19 and for coil 1 of configuration C3_R17 as an example of highly concave sections. Right: rendering of coil 1 of C3_R17 showing the concave section and of coil 1 of C4_R19 demonstrating that simple convex coils are achievable as well.

used throughout the process. In this section, an *a posteriori* study, meaning after all the optimisation steps, is performed to check the robustness of the many configurations already presented previously. The method is very similar to the generation of perturbed samples in the optimisation; however, here, the coils that are perturbed are of finite dimension with nine filaments. Each filament of a single coil is perturbed in the same way to simulate deviations of the whole winding pack. The perturbations are done using systematic deviations that are present in coils of the same type and statistical errors that are different in every coil. Then, a QFM surface is calculated for each sample, and a VMEC equilibrium is generated to extract data on the perturbed quasisymmetry from it. This is performed for all samples, and for increasing perturbation amplitude σ, a one-sigma data point is only added when all the VMEC samples converge. It should be noted that σ represents the exponential pre-factor of the kernel in the GP along one dimension; in order for it to correspond to a three-dimensional perturbation, $\sqrt{3}$ should be factored in. Each sampling is performed with 128 samples. The results are shown in figure 9.

There are two visible categories: in the first one are configurations C3_R18, C3_R19 and C4_R18, for which the departure from quasisymmetry increases relatively fast with increasing perturbation amplitude compared with the rest of the configurations, which appear to be robust to any deviations up to more than 1.3 mm for the C4_R16 stellarator, without any significant deterioration of the QS quality. The case C3_R17 appears as an outlier because although its QS error is low as well as its relative increase, for deviations higher than 3.5 mm, it is impossible to get all the samples to converge, indicating that this configuration might find itself in a sharp local minima. This is possible since this is the only configuration where high curvature is allowed, thereby minimising the squared flux as much as possible. According to figure 9, this results in a configuration that is highly accurate but not robust. This, together with configuration C3_R16, is one of the two configurations that show that the squared flux alone is not enough to quantify the robustness of a stellarator, as it is intuitive to think that a highly accurate stellarator can allow more damage to its coils as it compensates for the loss of accuracy. Moreover, the data displayed in figure 9 indicate that there is an inherent robustness of a stellarator.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 19 Nov 2026 at 05:20:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0022377826102256

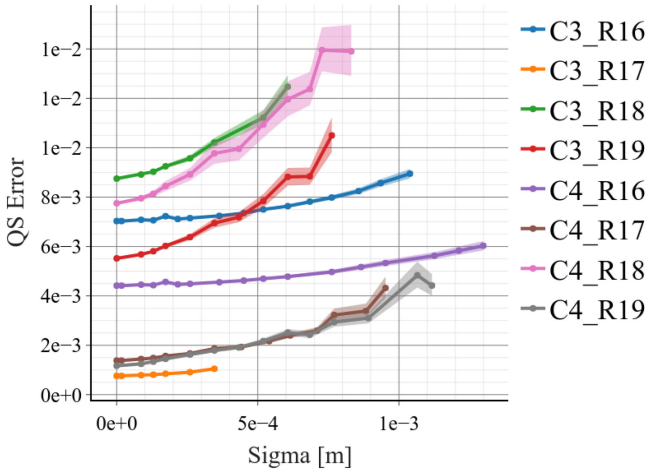


FIGURE 9. Robustness of the quasisymmetry versus coil perturbation amplitude for all the configurations. Here, the metric is the average deviation from quasisymmetry at a certain perturbation amplitude of the coils.

It is unclear at this stage what the decisive factor behind it is. An overview of the critical parameters of all the candidates presented here is shown under the form of star plots in Appendix E.

4.2. *Refining the optimisation*

Finally, in this section, the method for obtaining the best EPOS configuration, C4_R19, is described. This method is applied to other configurations, such as C4_R16 and C4_R18; however, it works with varying degrees of success. This technique is an additional process that follows the stochastic stage II, as presented in figure 10. The primary purpose of this method is to tailor the geometric characteristics of individual coils that do not present suitable parameters, such as HTS strain, curvature, size or inter-coil distance. For example, the former C4_R19 configuration has WL coils that are too big to fit in the vacuum chamber, or the former C4_R18 has coils 4–8 that possess concave sections. The overall improvement of the squared flux and quasisymmetry is shown in figure 10 for configuration C4_R19. The first step corresponds to one optimisation where the surface remains unchanged, and circular coils attempt to match it; therefore, QS stays constant, and the squared flux is considerably improved. Then, in the blue section, comes the stochastic single-stage phase, where a trade-off exists between the quality of the quasisymmetry and the accuracy of the coils. After this phase, the equilibria are no longer optimised, and the QS error remains constant. Following this, what is referenced here as stage III corresponds to a stochastic stage II, where only the coils are optimised. This yields a visible improvement in the squared flux, going from 1.8e−4 to 4e−6 m^2. It should be noted that this step also allows for all the coils to reach the desirable normal curvature everywhere. It requires numerous minor adjustments to the weights, specifically to the coil length penalty and mean-squared curvature. However, other metrics, such as HTS strain and the size of the WLs, are not satisfactory. That is why the final stage IV (in green) corresponds to a phase of adaptive and targeted strategies to correct imperfections. The first step, where a slight increase in the squared flux is

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 19 Nov 2026 at 05:20:14, subject to the Cambridge Core terms of use , available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0022377826102256

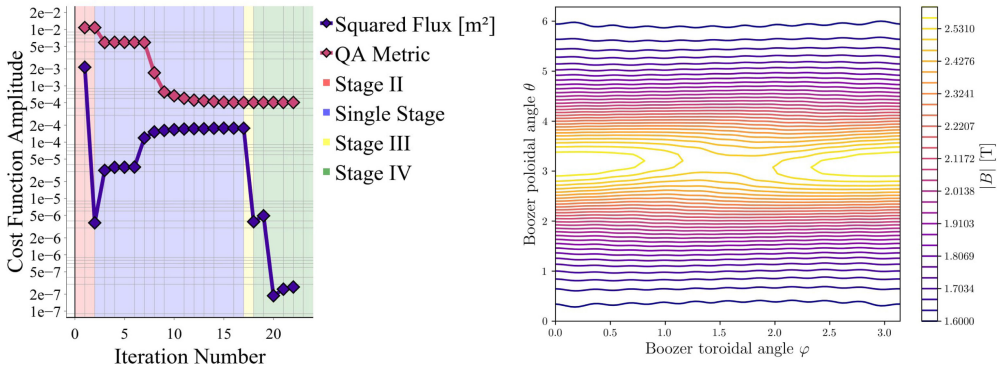


FIGURE 10. Left: amplitude of two cost functions during the different optimisation stages for C4_R19: the squared flux and the departure from quasisymmetry. Right: Boozer plot at the LCFS for the C4_R19 configuration obtained from a VMEC equilibrium after optimising a QFM surface, note that the field maximum at 2.6 T corresponds to the amplitude on the inboard side of the stellarator.

visible, refers to the fixing of all the degrees of freedom of the coils, except for those of the WL coils, together with an increase in the coil length penalty, and results in a successful decrease in the size of the WLs to fit in the vacuum chamber.

Moreover, in order to address the loss in squared flux from the previous stage, an additional correction step is performed where the degrees of freedom of the WLs are fixed (preventing them from expanding further) and all the other ones are freed, including the current degrees of freedom of all the coils. This is done as the WLs alone do not allow for considerable improvement of the squared flux (only by a maximal increment of 5e−7 m^2) compared with the 10 other coils in the field period. Figure 10 shows the increase in field accuracy yielded by this correction from 5e−6 to 1.8e−7 m^2. This enhancement by an order of magnitude is achieved while conserving the normal curvature to desirable degrees, improving the HTS strain levels, and decreasing the current ratio from 4 to 1.8. All the configurations are optimised with identical conditions. The case C4_R19 is not optimised with the intention of going to a ratio of 1.8. Through the manipulation of the hyperparameters and the specific local landscape of the parameter space, it naturally relaxed to such a current ratio, which does not happen to the other configurations. The last two steps, which slightly deteriorate the squared flux, are taken into account to compensate for the coil-to-coil distance being too small to accommodate the support material. On the right of the optimisation procedure in figure 10 is shown the Boozer plot for the LCFS of the device, displaying well-defined straight lines. The effects of coil ripple are also visible in the small-wavelength perturbations of the lines. The result is one configuration that fulfils all the preliminary requirements for a buildable pair plasma stellarator with HTS; the configuration is shown in figure 12. Planar-like WL coils are a desirable feature, as they simplify the winding procedure, especially for configurations with a current ratio of 1.8, which require only 1.8 times the number of windings of HTS tape as other coils. Furthermore, the WLs are optimised to leave enough space between the inboard side and the WLs as J × B force calculations (see Appendix B) show that if they remain too close to the standard coils, this can create areas where the forces can become problematic for the integrity of the support structure. It should be noted that one advantage for the optimiser of making

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 19 Nov 2026 at 05:20:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0022377826102256

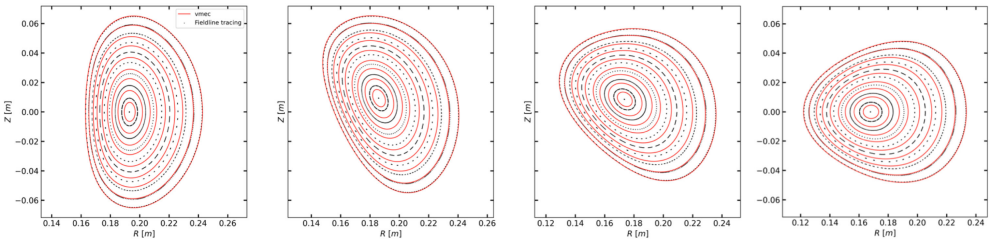


FIGURE 11. Poincaré cross-sections of the C4_R19 configuration at four different toroidal positions, comparison between the VMEC solution (red) and field line tracing (black) from the magnetic field generated by coils for different toroidal cross-sections, $\phi = [0, \pi/6, \pi/3, \pi/2, 5\pi/6, \pi]$. No island chains are visible.

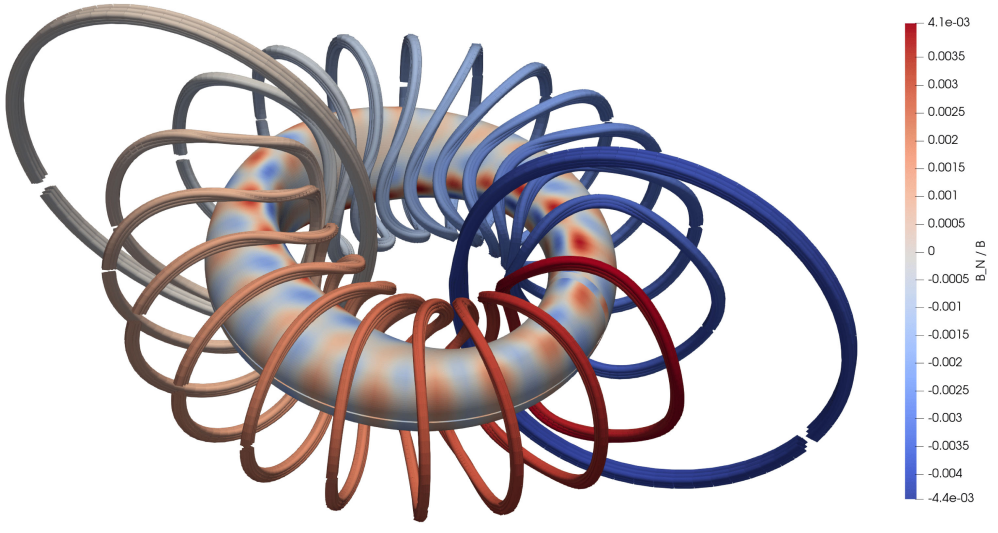


FIGURE 12. The EPOS C4_R19 configuration with coils and LCFS, the heatmap corresponds to the field error $\langle B \cdot n \rangle / \langle B \rangle$. The coils fulfil the engineering requirements and the equilibrium fulfils the physics requirements.

large WL coils and additionally freeing the coil current degrees of freedom is that it becomes easier to minimise the optimisation function by compressing all the other coils together (covering any gap previously due to the high currents of the WLs) and searching for paths to shape the standard coils as if EPOS were a regular stellarator symmetric device, which leads to simpler and repeated coils.

The extent of refinement observed in the C4_R19 configuration is not seen in other configurations where the same technique is applied. However, other geometric issues could be resolved, as mentioned previously. This shows that targeting localised issues of the configuration works as a straightforward method to exit local minima and access other solutions. Fixing specific degrees of freedom acts as a restriction of the parameter space, which facilitates the search for the optimiser. Additionally, it should be noted that the primary factor driving the continuous adjustment of the weights and necessitating ongoing refinement is the curvature of the coils. In most cases, further improvement of the squared flux requires a trade-off with the

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 19 Nov 2026 at 05:20:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0022377826102256

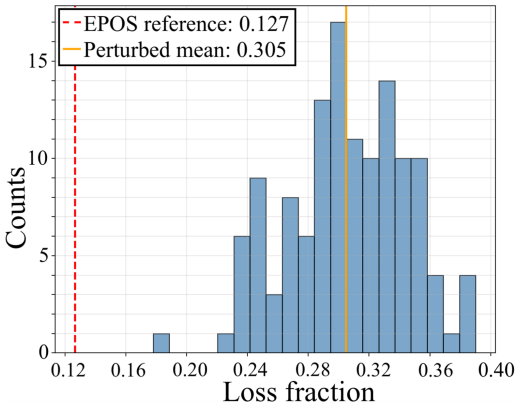


FIGURE 13. Loss fraction distribution for 128 perturbed EPOS C4_R19 configurations. Each count represents one EPOS stellarator where the coils have been randomly perturbed with 1 mm deviations. A thousand particles in the simulation are followed for 0.5 s with 1 eV each starting at the radial coordinate of $s = 0.9$.

concavity of the coils. However, this is not the case for the C4_R19 configuration. One hypothesis explaining the available minimum is that the equilibrium LCFS is convex and does not exhibit significant shaping, allowing for simple coils to be found. The Poincaré cross-sections of this configuration are shown in figure 11. No visible island chains are present; the flux surfaces do not possess considerable shaping, as the average iota is 0.06.

An additional check of how manufacturing and assembly deviations directly affect particle confinement is presented in figure 13. The distribution of the loss fraction of particles is shown after following 1000 trajectories for 0.5 s with SIMPLE. The particles carry an energy of 1 eV and are initialised on the outboard side of the stellarator at a radial coordinate of $s = 0.9$. The distribution is compared with the unperturbed C4_R19 configuration which loses 12.7 % of the particles. Perturbing the coils with 1 mm deviations deteriorates the confinement and results in an average loss of 30 % of electrons and positrons reaching in some cases 38 %. These results do not account for two experimental factors that work in favour of better confinement: the velocity distribution used in these simulations is random, however, the positron beam feeding the experiment is expected to have on average a greater parallel velocity than perpendicular. Then, the simulations do not account for cyclotron cooling that radiates a considerable amount of the perpendicular energy. Given the dimensions of the stellarator, having 30 % particle loss after 0.5 s with coil perturbation amplitudes of 1 mm, with random perpendicular velocities and no cyclotron cooling, is a promising result.

Further post optimisation studies regarding the quality of the C4_R19 configuration are also performed through large-scale scans where 12,600 different EPOS stellarators are considered (see Appendix C). These scans enable us to compare, at least locally, the results obtained so far. The primary coil characteristics of C4_R19 seem to outperform all the configurations from the scans. Although the QS error of this configuration is higher than most configurations found in the scans, what makes it stand out particularly is the field accuracy that can be achieved with relatively simple coils.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 19 Nov 2026 at 05:20:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0022377826102256

5. Conclusion

In this paper, the procedure to find suitable candidates for the EPOS stellarator is presented. With the ultimate goal of confining pair plasmas, the final configuration is able to fulfil a series of physics and engineering constraints that restrict the parameter space of possible buildable stellarators. A determining factor for EPOS is the size requirement needed to achieve densities high enough to observe collective behaviour. Concretely, this is expressed in the search for configurations that possess low iota for coils with low HTS strain, no concavity, WLs with stray fields for the injection of particles, sufficient coil-to-coil distance to accommodate material between them and improved robustness to manufacturing errors of the coils. These conditions led to the development of a minimisation scheme that includes stochastic optimisation, as well as single-stage methods and stage II techniques with finite-build coils, HTS strain and a subsequent restriction of the parameter space.

To assess different configurations, a scan is performed at the major radius and in the current ratio between the standard coils and the WLs. Although the major radii vary only between 16 and 19 cm and the current ratios between 3 and 4, considerable differences are visible among all the configurations, providing an indication of the breadth of the optimisation space for three-dimensional configurations. Single particle guiding-centre orbit simulations showed that confinement times of approximately 2 s with losses inferior to 10 % are possible when starting the particles at the magnetic axis. Moreover, *a posteriori* robustness verification for the configurations not only showed that some are tolerant to perturbations exceeding 1 mm but also revealed an indication of intrinsic robustness in several configurations. Finally, all the candidates achieved below threshold values for the HTS strain, and all those with a current ratio of 4 found convex shaped coils. Specifically, EPOS C4_R19 displayed good characteristics in all the metrics and will be studied further. This study presents a possible approach to optimising the design of a stellarator, enabling its complete construction. Following this design stage, a subsequent engineering phase is being conducted to assess the feasibility of the configuration (Huslage *et al*. 2025; Gil *et al*. 2026*a*). This includes analysis of the stresses on the coils, the magnetic field permeability of the materials used in construction and even positron injection under realistic conditions to enable a comprehensive design of the coils.

Looking towards future devices, the integration of engineering proxies is highly valuable as it can guide optimisation, taking into account, for example, von Mises/Tresca stresses instead of solely forces, field errors from magnetic permeability of materials and the current distribution in the winding packs, etc.

Finally, if EPOS achieves its density targets, due to average ‘bad curvature’, there is a possibility that the electrostatic version of interchange modes can still arise. This, however, would not be catastrophic, quite the opposite in fact, as it would be the first time MHD instabilities were observed in a matter–antimatter plasma.

Acknowledgements

We thank M. Landreman, J. Loizu, T.S. Pedersen, E. Paul, J. von der Linden and A. Deller for useful discussions and insights.

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

Funding

This work was supported through grants from the Helmholtz Association Young Investigators Group program as project VH-NG-1430.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 19 Nov 2026 at 05:20:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0022377826102256

Declaration of interests

The authors report no conflict of interest.

Appendix A. Number of coils

To determine the number of coils in the configurations studied above, a study is conducted on the quality of quasisymmetry as a function of the number of coils per half-field period and the distance between coils and the LCFS.

The REGCOIL optimisation software (Landreman 2017) is employed. REGCOIL solves a convex optimisation problem to find the current potential that best reproduces the target plasma surface on a predetermined coil winding surface. One of the advantages of this procedure is that it has only one global minimum, thereby allowing for accurate comparisons between different coil-to-surface distances, for example. Once the current potential is found, the software is then able to 'cut coils' from the winding surface, and here is where the number of coils is set by the user. Then, once a configuration with coils is found, a check of the quality of the quasisymmetry is performed using the previously mentioned QFM method. Figure 14 summarises the findings for this EPOS configuration.

As expected, the QS error drops as the number of coils increases. This can be explained by the reduction of coil ripple due to the increasing coverage of the coils. However, after seven coils, it is visible that there is a saturation in the quality of quasisymmetry, which is observed across all coil-to-surface distances. Moreover, if the 2 cm coil-to-surface distance is not taken into account, saturation is even achieved at six coils per half-field period. Since six coils per half-field period with finite dimensions are too cumbersome for the device, especially at 16 cm in major radius, it is decided to set the number of coils at 11 coils per field period, including WL coils.

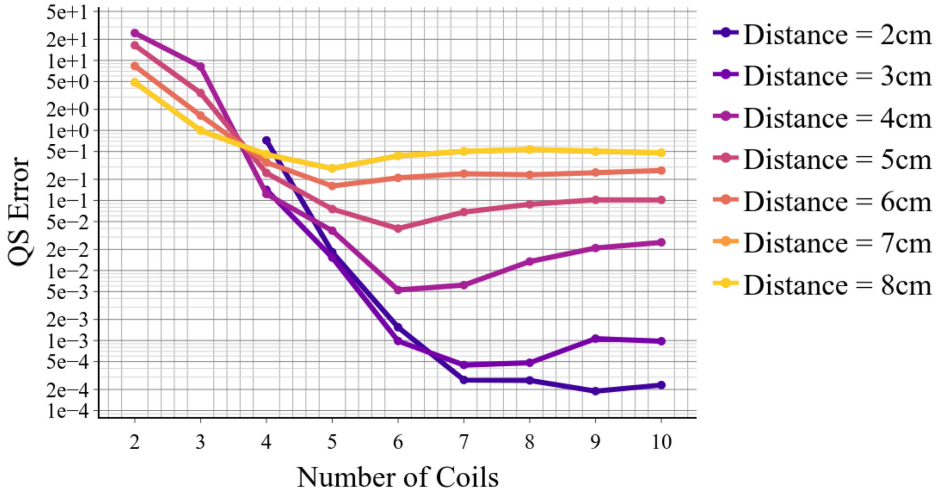


FIGURE 14. Quasisymmetric error as a function of the number of coils and of the coil to surface distance.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 19 Nov 2026 at 05:20:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0022377826102256

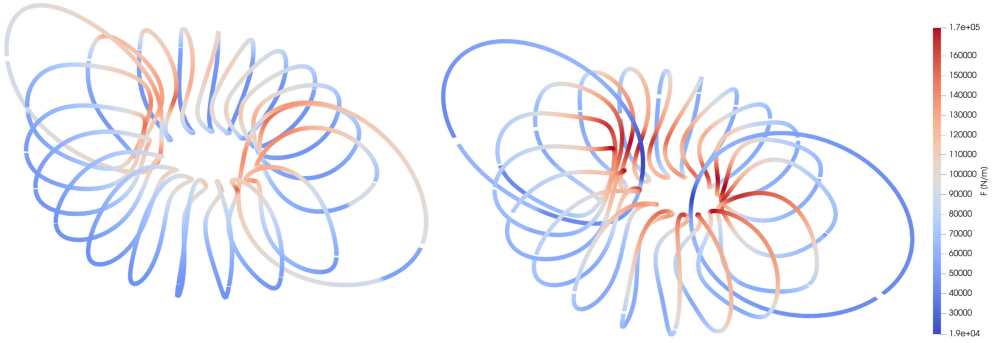


FIGURE 15. Simulated coil forces shown on the coil filaments expressed in Nm^{-1} on two EPOS candidates: C3_R16 (left) and C4_R19 (right), spatial dimensions are to scale, apparent discontinuities in the left plot are due to applying a common scale for both stellarators.

Appendix B. Coil forces

Given the small dimensions and high fields of the EPOS device, mechanical stresses due to J × B forces are checked in this section. Using the force calculation framework described in (Landreman, Hurwitz & Antonsen 2025) and implemented in SIMSOPT (Hurwitz, Landreman & Kaptanoglu 2024), a three-dimensional force vector in each of the quadrature points of each coil is calculated and estimated in in Nm^{-1}. This vector includes contributions from forces due to the interaction between the current in a given coil and its own generated field (self-field), as well as the external field induced by the remaining coils. These calculations are obtained by assuming a coil 'wire' with a finite rectangular cross section of 6 mm × 12 mm, which is close to the expected experimental values of EPOS. In figure 15, the amplitude of the forces for two different candidates, C3_R16 (left) and C4_R19, the latter being the best candidate so far, is displayed.

These two configurations show possible differences in feasibility between two design choices: C3_R16 has a smaller major radius and coils (when compared with C4_R19), which can lead to smaller coil-to-coil distances and higher forces. However, because of the 2 T requirement on the axis, this also means that larger coils require larger currents. Combined, this leads to comparable amplitudes between the two configurations, with the maximum value for C3_R16 being 1.5e5 Nm^{-1} and for C4_R19 1.7e5 Nm^{-1}. The overall higher force values for the C4_R19 are observed in the inboard side, which is expected since the coil-to-coil distance decreases in that region. One additional feature of the larger configuration is the spacing between the inboard side of the standard coils and the WL coils. Although most coils experience increased stresses, the WL coils exhibit considerably reduced forces compared with their smaller counterparts. This is mainly due to two factors: first, as mentioned previously, C4_R19 only possesses a factor of 1.8 between the currents of the standard coils and the WL coils, and unlike C3_R16 on the inboard side, the WL coils are further apart from the standard coils.

Appendix C. Scans of the configuration space

Following the 8 EPOS configuration study presented in previous sections, where the configurations are extensively optimised to present viable candidates, large-scale

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 19 Nov 2026 at 05:20:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0022377826102256

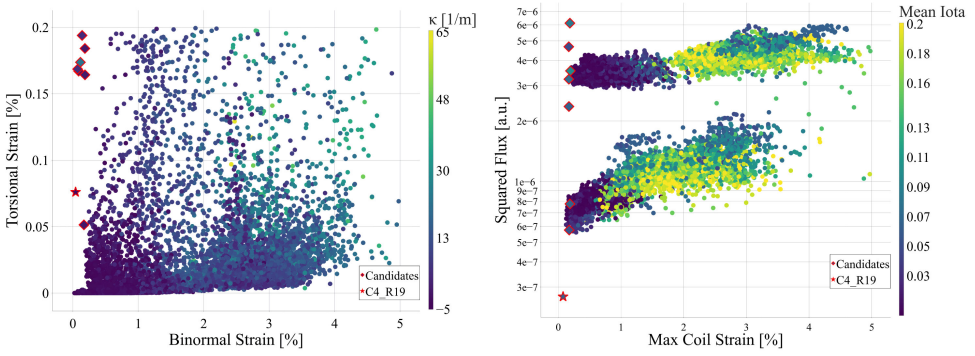


FIGURE 16. Left: scatter plot showing the torsional strain versus the binormal strain, where the intensity shows the convexity of the coils of the EPOS large-scale scans. Right: squared flux versus the maximum measured coil strain of a configuration defined as the maximum of the binormal and the torsional strain together with the mean iota of each configuration as a heatmap.

scans allow us to observe where these candidates exist in the N-dimensional stellarator configuration space. The principle of the scans presented here is to target large numbers of possible EPOS configurations without requiring that they be buildable, in order to sample from different regions in the parameter space, possibly providing unexplored better configurations. Note that, unlike the optimisation explained in this work, the scans are performed following a traditional stage-I/stage-II method.

Main target parameters for the EPOS equilibria are: major radii between 16 and 19 cm, iota between 0.05 and 0.20 and aspect ratio between 3 and 6, the magnetic well is also varied by varying the magnitude of its respective hyperparameter, as for the coils the ratio between the WL coils and the standard coils is scanned between 2 and 4, and the coil length is scanned through its hyperparameter variation. The total number of different coil sets found is 12 600, with effectively 2100 different EPOS equilibria, and six different coil sets for each equilibrium. Note that these numbers are still low compared with the uncountable measure of the parameter space. Nonetheless, the information provided by these scans allows us to place the configurations found in this work at least locally. As observed in figure 16, the left plot shows the three most important quantities regarding the coils: torsional versus binormal strain and coil convexity κ. With the red outlines, the data points referring to the fully optimised configurations are plotted. Compared with the 12 600 data points they display optimal strains and are located in a region unreached by the scans at very low binormal strains ($\epsilon \leqslant 0.2\,\%$) and convex coils ($\kappa \leqslant 0$). Moreover, in the right plot, where the quadratic flux is plotted against the maximum coil strain, and considering the rotational transform of the configurations, these configurations appear again at low strain, with similar field accuracy, but at a higher iota. Note that the C4_R19 configuration stands out in all of the parameters shown here.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 19 Nov 2026 at 05:20:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0022377826102256

Appendix D. Parameters of the C4_R19 configuration

See table 2.

Parameters	C4_R19
Max $B \cdot n/B$	4.1e−3
$\langle B \cdot n \rangle / \langle B \rangle$	9.6e−4
QS error [a.u]	1.1e−3
QS at $\sqrt{3}\sigma = 1$ mm [a.u.]	4e−3
Minor radius [cm]	5.12
Major radius [cm]	19
Aspect ratio	3.71
Mean rotational transform ι	0.064
Well [%]	−21.3
Required positrons at $T_e = 1$ eV	∼1e10
a/λ_D at $T_e = 1$ eV	9.73
Equilibrium volume [l]	9.87
Avg coil length (Std/WL) [cm]	86/154
Total coil length [m]	9.53
Min coil-to-plasma distance [cm]	∼3.7*
Min. coil-to-coil distance [cm]	∼1.5*
Max. force on conductor [N m^{-1}]	1.7e5
Currents (Std/WL) [kA]	77.7/142.7
Max binormal strain [%]	0.041
Max torsional strain [%]	0.076
Max concavity κ [m^{-1}]	−2.57
Max. curvature [m^{-1}]	18.9
Max. MSC curvature [m^{-1}]	103.57

TABLE 2. Main parameters of the C4_R19 configuration of EPOS, standing as the best candidate so far. *Note that coil-to-surface and coil-to-coil distance values consider already SIMSOPT finite-build modifications; however, these are still subject to variations depending on the final engineering design that accounts for finite CAD dimensions.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 19 Nov 2026 at 05:20:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0022377826102256

Appendix E. Overview of the EPOS candidates configurations

See figures 17 and 18.

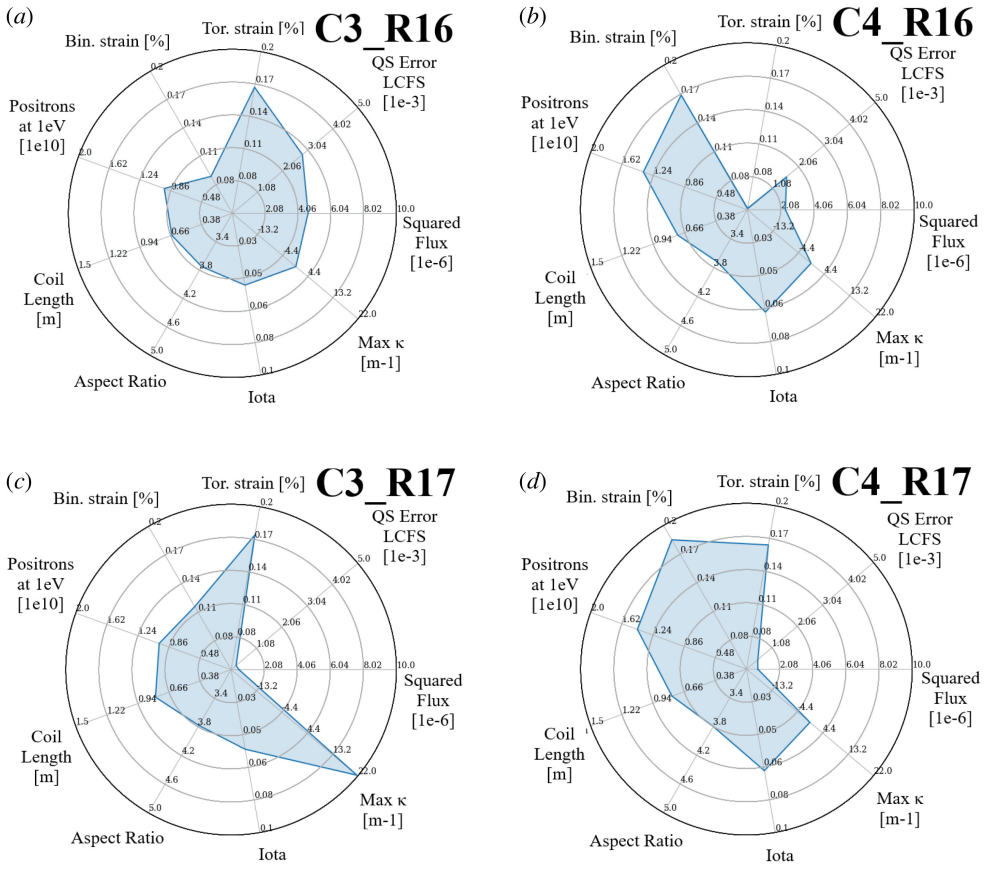


FIGURE 17. Star plots summarising the relevant properties for the EPOS stellarator candidates with major radii 16 and 17 cm. For all the properties except for iota, the blue area should ideally be minimised.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 19 Nov 2026 at 05:20:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0022377826102256

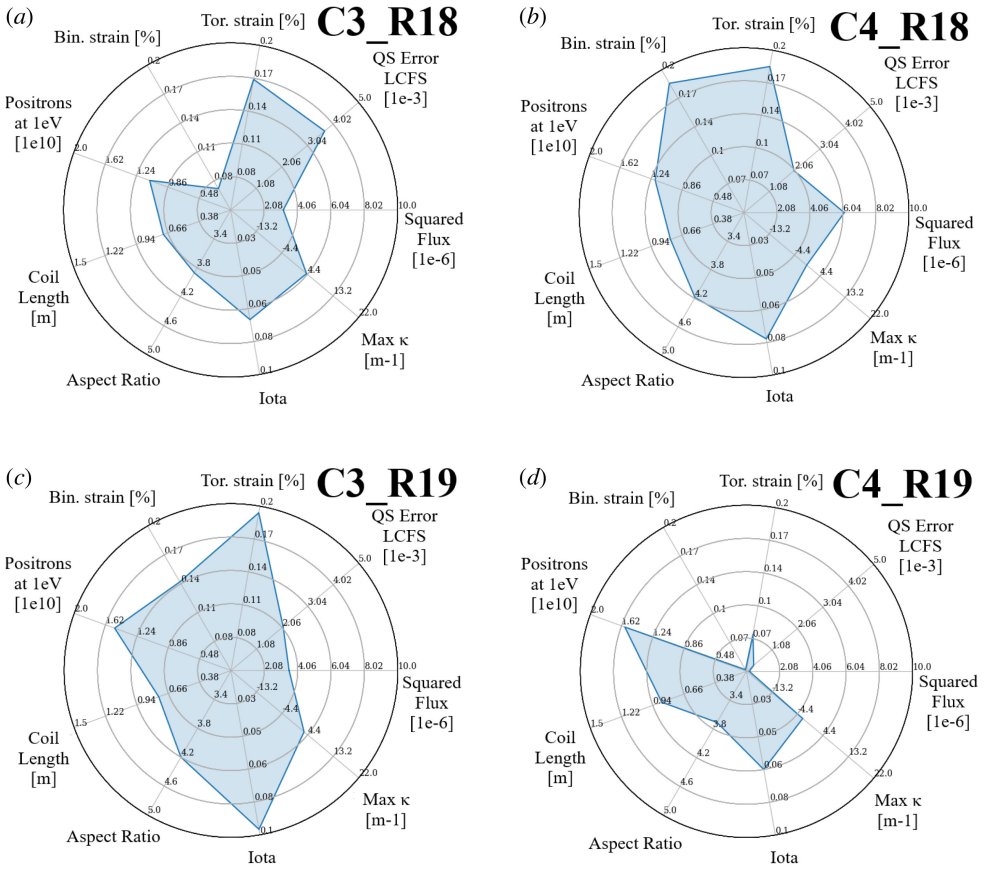


FIGURE 18. Star plots summarising the relevant properties for the EPOS stellarator candidates with major radii 18 and 19 cm. For all the properties except for iota, the blue area should ideally be minimised.

REFERENCES

ALBERT, C.G., KASILOV, S.V. & KERNBICHLER, W. 2020*a* Accelerated methods for direct computation of fusion alpha particle losses within, stellarator optimization. *J. Plasma Phys.* **86**, 815860201.

ALBERT, C.G., KASILOV, S.V. & KERNBICHLER, W. 2020*b* Symplectic integration with non-canonical quadrature for guiding-center orbits in magnetic confinement devices. *J. Comput. Phys.* **403**, 109065.

ANDEREGG, F., HOLLMANN, E.M. & DRISCOLL, C.F. 1998 Rotating field confinement of pure electron plasmas using trivelpiece-gould modes. *Phys. Rev. Lett.* **81**, 4875–4878.

ANDREEVA, T., BRÄUER, T., ENDLER, M., KISSLINGER, J. & IGITKHANOV, Y. 2004 Analysis of the magnetic field perturbations during the assembly of wendelstein 7-x. *Fusion Sci. Technol.* **46**, 388–394.

ARONS, J. 1979 Some problems of pulsar physics or i'm madly in love with electricity. *Space Sci. Rev.* **24**, 437–510.

BAILLOD, A., PAUL, E.J., RAWLINSON, G., HAQUE, M., FREIBERGER, S.W. & THAPA, S. 2025 Integrating novel stellarator single-stage optimization algorithms to design the columbia stellarator experiment. *Nucl. Fusion* **65**, 026046.

BAILLOD, A., VEKSLER, A., LOPEZ, R., SCHMELING, D., CAMPAGNA, M., PAUL, E. & KNYAZEV, A. 2026 Update on the design of the columbia stellarator experiment. arXiv: 2601.00673.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 19 Nov 2026 at 05:20:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0022377826102256

GIL, P.F., BRACK, V., SCHULER, T., HUSLAGE, P., SMONIEWSKI, J. & STENSON, E.V. 2026*a* M*a*nufacturing tolerances of non-planar coils for an optimized tabletop stellarator. *IEEE Trans. Appl. Supercond.* **36**, 1–5.

GIL, P.F., SMONIEWSKI, J., JORGE, R., HUSLAGE, P. & STENSON, E.V. 2026*b* Stochastic single-stage stellarator optimization using fixed-boundary equilibria. arXiv: 2603.11699.

GIULIANI, A., WECHSUNG, F., CERFON, A., LANDREMAN, M. & STADLER, G. 2023 Direct stellarator coil optimization for nested magnetic surfaces with precise quasi-symmetry. *Phys. Plasmas* **30**, 042511.

GIULIANI, A., WECHSUNG, F., STADLER, G., CERFON, A. & LANDREMAN, M. 2022 Direct computation of magnetic surfaces in boozer coordinates and coil optimization for quasisymmetry. *J. Plasma Phys.* **88**, 905880401.

HAHN, S., PARK, D.K., BASCUNAN, J. & IWASA, Y. 2011 Hts pancake coils without turn-to-turn insulation. *IEEE Trans. Appl. Supercond.* **21**, 1592–1595.

HAHN, S., RADCLIFF, K., KIM, K., KIM, S., HU, X., KIM, K., ABRAIMOV, D.V. & JAROSZYNSKI, J. 2016 'Defect-irrelevant' behavior of a no-insulation pancake coil wound with rebco tapes containing multiple defects. *Supercond. Sci. Technol.* **29**, 105017.

HEGNA, C.C. & NAKAJIMA, N. 1998 On the stability of Mercier and ballooning modes in stellarator configurations. *Phys. Plasmas* **5**, 1336–1344.

HELANDER, P. 2014 Microstability of magnetically confined electron-positron plasmas. *Phys. Rev. Lett.* **113**, 135003.

HELANDER, P. & CONNOR, J.W. 2016 Gyrokinetic stability theory of electron–positron plasmas. *J. Plasma Phys.* **82**, 905820301.

HENNEBERG, S.A., HUDSON, S.R., PFEFFERLÉ, D. & HELANDER, P. 2021 Combined plasma–coil optimization algorithms. *J. Plasma Phys.* **87**, 905870226.

HIRSHMAN, S.P. & WHITSON, J.C. 1983 Steepest-descent moment method for three-dimensional magnetohydrodynamic equilibria. *Phys. Fluids* **26**, 3553–3568.

HUDSON, S.R., ZHU, C., PFEFFERLÉ, D. & GUNDERSON, L. 2018 Differentiating the shape of stellarator coils with respect to the plasma boundary. *Phys. Lett. A* **382**, 2732–2737.

HURWITZ, S., LANDREMAN, M. & KAPTANOGLU, A. 2024 Electromagnetic coil optimization for reduced lorentz forces. arXiv: 2410.09337.

HUSLAGE, P., KULLA, D., LOBSIEN, J.-F., SCHULER, T. & STENSON, E.V. 2024*a* Winding *a*ngle optimization and testing of small-scale, non-planar, high-temperature superconducting stellarator coils. *Supercond. Sci. Technol.* **37**, 085010.

HUSLAGE, P., PAUL, E.J., GIL, M.H.P.F., FOPPIANI, N., SMONIEWSK, J. & STENSON, E.V. 2024*b* Strain optimization for rebco high-temperature superconducting stellarator coils in simsopt. arXiv: 2409.01925.

HUSLAGE, P., SCHULER, T., GIL, P.F., BRACK, V., SCHMELING, D., ORONA, D.A.R., SCHOENBERG, E. VON, THUN, T., SMONIEWSKI, J. & STENSON, E.V. 2025 A non-planar rebco test coil with 3d-printed aluminum support structure for the epos stellarator. *IEEE Trans. Appl. Supercond.* **36**, 1–6.

IMBERT-GÉRARD, L.-M., PAUL, E.J. & WRIGHT, A.M. 2024 *An Introduction to Stellarators: From Magnetic Fields to Symmetries and Optimization*. Society for Industrial and Applied Mathematics,

JORGE, R., GOODMAN, A., LANDREMAN, M., RODRIGUES, J. & WECHSUNG, F. 2023 Single-stage stellarator optimization: combining coils with fixed boundary equilibria. *Plasma Phys. Control. Fusion* **65**, 074003.

KENNEDY, D. & HELANDER, P. 2021 Coulomb collisions in strongly anisotropic plasmas I. cyclotron cooling in electron–ion plasmas. *J. Plasma Phys.* **87**, 905870107.

KIM, G. *et al.* 2025 Numerical study of screening current on hts stellarator coil and experimental validation through screening current induced voltage. *IEEE Trans. Appl. Supercond.* **35**, 1–5.

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

LANDREMAN, M., HURWITZ, S. & ANTONSEN, T.M. 2025 Efficient calculation of self magnetic field, self-force, and self-inductance for electromagnetic coils with rectangular cross-section. *Nucl. Fusion* **65**, 036008.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 19 Nov 2026 at 05:20:14, subject to the Cambridge Core terms of use , available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0022377826102256

LANDREMAN, M., MEDASANI, B., WECHSUNG, F., GIULIANI, A., JORGE, R. & ZHU, C. 2021 SIMSOPT: a flexible framework for stellarator optimization. *J. Open Source Softw.* **6**, 3525.

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. [Computer Software]. Available at: https://doi.org/10.11578/dc.20180627.6.

LOBSIEN, J.-F., DREVLAK, M., PEDERSEN, T.S. & Team, W7-X 2018 Stellarator coil optimization towards higher engineering tolerances. *Nucl. Fusion* **58**, 106013.

PAZ-SOLDAN, C. 2020 Non-planar coil winding angle optimization for compatibility with non-insulated high-temperature superconducting magnets. *J. Plasma Phys.* **86**, 815860501.

PEDERSEN, T.S. *et al.* 2016 Confirmation of the topology of the Wendelstein 7-X magnetic field to better than 1:100 000. *Nat. Commun.* **7**, 13493.

PEDERSEN, T.S., DANIELSON, J.R., HUGENSCHMIDT, C., MARX, G., SARASOLA, X., SCHAUER, F., SCHWEIKHARD, L., SURKO, C.M. & WINKLER, E. 2012 Plans for the creation and studies of electron–positron plasmas in a stellarator. *New J. Phys.* **14**, 035010.

QIAN, T.M. *et al.* 2023 Design and construction of the muse permanent magnet stellarator. *J. Plasma Phys.* **89**, 955890502.

RODRÍGUEZ, E., PAUL, E.J. & BHATTACHARJEE, A. 2022 Measures of quasisymmetry for stellarators. *J. Plasma Phys.* **88**, 905880109.

SINGH, L., KRUGER, T.G., BADER, A., ZHU, C., HUDSON, S.R. & ANDERSON, D.T. 2020 Optimization of finite-build stellarator coils. *J. Plasma Phys.* **86**, 905860404.

STENSON, E.V. *et al.* 2018 Lossless positron injection into a magnetic dipole trap. *Phys. Rev. Lett.* **121**, 235005.

STRICKLER, D.J., BERRY, L.A. & HIRSHMAN, S.P. 2002 Designing coils for compact stellarators. *Fusion Sci. Technol.* **41**, 107–115.

STRYKOWSKY, R.L., BROWN, T., CHRZANOWSKI, J., COLE, M., HEITZENROEDER, P., NEILSON, G.H., REJ, D. & VIOL, M. 2009 Engineering cost and schedule lessons learned on ncsx. In *2009 23rd IEEE/NPSS Symposium on Fusion Engineering*, pp. 1–4. IEEE (Institute of Electrical and Electronics Engineers).

TSYTOVICH, V. & WHARTON, C.B. 1978 Laboratory electron-positron plasma - a new research object. *Comments Plasma Phys. Control. Fusion* **4**, 91–100.

WANG, Y., CHAN, W.K. & SCHWARTZ, J. 2016 Self-protection mechanisms in no-insulation (re)ba2cu3ox high temperature superconductor pancake coils. *Supercond. Sci. Technol.* **29**, 045007.

WECHSUNG, F., GIULIANI, A., LANDREMAN, M., CERFON, A. & STADLER, G. 2022*a* Single-stage gradient-based stellarator coil design: stochastic optimization. *Nucl. Fusion* **62**, 076034.

WECHSUNG, F., GIULIANI, A., LANDREMAN, M., CERFON, A. & STADLER, G. 2022*b* Stochastic and a posteriori optimization to mitigate coil manufacturing errors in stellarator design. *Plasma Phys. Control. Fusion* **64**, 105021.

ZHU, C., HUDSON, S.R., SONG, Y. & WAN, Y. 2017 New method to design stellarator coils without the winding surface. *Nucl. Fusion* **58**, 016008.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 19 Nov 2026 at 05:20:14, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0022377826102256