

A family of quasi-axisymmetric stellarators with varied rotational transform

S. Buller^{1,2,†}, M. Landreman¹, J. Kappel¹ and R. Gaur^{1,3}

¹Institute for Research in Electronics and Applied Physics, University of Maryland, MD 20742, USA

²Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08543, USA

³Department of Mechanical and Aerospace Engineering, Princeton University, Princeton, NJ 08543, USA

(Received 17 January 2024; revised 27 September 2024; accepted 1 October 2024)

We apply a continuation method to recently optimized stellarator equilibria with excellent quasi-axisymmetry to generate new equilibria with a wide range of rotational transform profiles. Using these equilibria, we investigate how the rotational transform affects fast-particle confinement, the maximum coil–plasma distance, the maximum growth rate in linear gyrokinetic ion-temperature gradient simulations and the ion heat flux in corresponding nonlinear simulations. We find values of two-term quasi-symmetry error comparable to or lower than those of the similar Landreman–Paul (*Phys. Rev. Lett.*, vol. 128, 2022, 035001) configuration for values of the mean rotational transform $\bar{\iota}$ between 0.12 and 0.75. The fast-particle confinement improves with $\bar{\iota}$ until $\bar{\iota} = 0.73$, at which point the degradation in quasi-symmetry outweighs the benefits of further increasing $\bar{\iota}$. The required coil–plasma distance only varies by about $\pm 10\%$ for the configurations under consideration, and is between 2.8 and 3.3 m when the configuration is scaled up to reactor size (minor radius $a = 1.7$ m and volume-averaged magnetic field strength of 5.86 T). The maximum growth rate from linear gyrokinetic simulations increases with $\bar{\iota}$, but also shifts towards higher k_y values. The maximum linear growth rate is sensitive to the choice of flux tube at rational $\bar{\iota}$, but this can be compensated for by taking the maximum over several flux tubes. The corresponding ion heat fluxes from nonlinear simulations display a non-monotonic relation to $\bar{\iota}$. Sufficiently large positive shear is destabilizing. This is reflected in both linear growth rates and nonlinear heat fluxes.

Keywords: fusion plasma

1. Introduction

Stellarators are toroidal magnetic confinement devices that rely on non-axisymmetric magnetic fields to confine a plasma. Non-axisymmetry makes it possible to produce the confining magnetic field using currents external to the plasma. However, without an explicit symmetry direction, particles are not guaranteed to be confined by conservation laws, and the magnetic field must be optimized to confine particles.

† Email address for correspondence: sb0095@princeton.edu

One possible optimization strategy is to optimize for a property known as quasi-symmetry (Boozer 1983), in which the magnitude of the magnetic field has a symmetry in a special set of coordinates. Over the last few years, tremendous progress has been made in optimizing stellarator geometries to achieve excellent quasi-symmetry (Landreman & Paul 2022). However, the results so far are mostly limited to configurations scattered in parameter space, which makes it hard to systematically study the effects of varying equilibrium parameters, such as the rotational transform ι and global magnetic shear \hat{s} .

In this paper, we attempt to alleviate this issue by applying continuation methods to continuously deform the rotational transform profile of the ‘Precise QA’ vacuum configuration described in Landreman & Paul (2022). This procedure lets us generate a family of stellarator geometries that are similar to each other in a sense that is explained in the next section, but where the rotational transform profile varies.

Since the mean rotational transform of the original ‘Precise QA’ configuration was arbitrarily imposed, it is not *a priori* obvious that there are any inherent trade-offs between quasi-symmetry and different values of mean ι . Increasing the rotational transform is expected to improve the fast-particle confinement, since the width of banana orbits scales $\propto 1/\iota$ (Paul *et al.* 2022), and could thus be beneficial for a reactor.

Optimization by continuation is a technique whereby an optimization problem is solved by continuously deforming a simpler optimization problem. If each deformation is a small perturbation to the previous problem, the solution of the previous problem can be used as a good initial guess to the perturbed problem. By repeatedly applying small perturbations and optimizing at each step, we eventually arrive at a solution to the novel optimization problem.

A variety of continuation-type methods have been used for stellarator equilibrium and optimization calculations previously. The code DESC allows for magnetohydrodynamic equilibria to be computed by continuation in the plasma boundary shape or plasma profiles (Conlin *et al.* 2023). One common optimization approach – optimizing first in a small parameter space of Fourier modes describing the plasma shape, then expanding the number of Fourier modes in the parameter space and re-optimizing – can be viewed as a type of continuation method. This approach was discussed for instance in Landreman & Paul (2022). Continuation is also a standard method for exploring the Pareto front in multi-objective optimization, and has recently seen use in this context to explore the trade-offs between quasi-symmetry and aspect ratio (Bindel, Landreman & Padidar 2023).

Larger systematic explorations of the space of quasi-symmetric stellarators have previously been published by Landreman (2022), which relied on the near-axis expansion to directly construct optimized stellarator equilibria (Landreman, Sengupta & Plunk 2019; Plunk, Landreman & Helander 2019). In contrast, our geometries are solutions to the full equilibrium equation, with a realistic number of Fourier modes describing the boundary. Giuliani *et al.* (2023) recently made public a large database of quasi-axisymmetric (QA) stellarators with varying $\bar{\iota}$, which could be used to perform studies similar to this one.

While there is a rich mathematical literature on continuation methods (also known as homotopy methods), see for example Allgower & Georg (2003), the application of the method to practical optimization problems is often based more on intuition than on mathematical rigour (Mobahi & Fisher 2015). This is also the approach taken here.

The rest of the paper is organized as follows. In the next section, we present the objective function used in the optimization, and introduce the continuation method. In § 3, we apply this method to generate ‘Precise QA’-like configurations with varying mean rotational transform or shear. We analyse these configurations with respect to fast-particle confinement, gyrokinetic stability and required coil–plasma distance. We find that fast-particle confinement improves with higher $\bar{\iota}$, up to a point, while the maximum

growth rate in our gyrokinetic simulations increases and shifts towards higher k_y . Finally, we discuss the results in § 4. All the equilibria presented here can be found as supplemental material on Zenodo at <https://zenodo.org/doi/10.5281/zenodo.10521393> (see Buller 2024).

2. Theory

A magnetic field \mathbf{B} is quasi-symmetric with helicity (M, N) if its magnitude $B = |\mathbf{B}|$ can be written as

$$B = B(\psi, M\vartheta - N\zeta), \tag{2.1}$$

where ψ is a flux-surface label, which we take to be the toroidal flux divided by 2π , and ϑ and ζ are magnetic angles, which we take to be the Boozer angles. Compared with the generic case where $B = B(\psi, \vartheta, \zeta)$ depends on all angles independently, the Lagrangian of a charged particle travelling through a quasi-symmetric field has a symmetry direction when written in Boozer coordinates, which ensures that the charged particle is confined by the magnetic field (Boozer 1983; Nührenberg & Zille 1988; Rodriguez, Helander & Bhattacharjee 2020). The symmetry direction depends on M and N : fields with $N = 0$ are known as QA; fields with non-zero N and M are called quasi-helically symmetric.

As the Boozer angles ϑ and ζ themselves depend on \mathbf{B} , it is not known how to directly compute fields with quasi-symmetry, except for the case where the fields possess exact axisymmetry ($\mathbf{B} = \mathbf{B}(\psi, \vartheta)$). Using different asymptotic methods for finding quasi-symmetry typically yields overdetermined systems of equations beyond a certain order (Garren & Boozer 1991; Plunk & Helander 2018), which suggests that exact quasi-symmetry apart from true axisymmetry may be impossible to achieve. The approach taken historically is to instead optimize \mathbf{B} to achieve approximate quasi-symmetry, which can yield configurations that are good enough in practice. This is also the approach taken in this paper.

As in Landreman & Paul (2022), we take our objective function to be

$$f = f_{\text{QS}} + (A - A_*)^2 + (\bar{\iota} - \bar{\iota}_*)^2, \tag{2.2}$$

where A is the aspect ratio, $\bar{\iota}$ is the mean (over ψ) of the rotational transform and f_{QS} is the so-called two-term formulation of quasi-symmetry (Helander & Simakov 2008). Specifically

$$f_{\text{QS}} = \sum_{s_j} \left\langle \left(\frac{1}{B^3} [N - \iota M] \mathbf{B} \times \nabla B \cdot \nabla \psi - [MG + NI] \mathbf{B} \cdot \nabla B \right)^2 \right\rangle, \tag{2.3}$$

where the sum is over flux surfaces, $\langle \cdot \rangle$ denotes the flux-surface average and $2\pi G/\mu_0$ and $2\pi I/\mu_0$ are the poloidal current outside and toroidal current inside the flux surface, respectively. The f_{QS} objective is zero when \mathbf{B} is quasi-symmetric with helicity (M, N) , so different kinds of quasi-symmetry can be targeted by specifying M and N . The second and third terms on the right-hand side of (2.2) target a specific aspect ratio A_* and mean rotational transform $\bar{\iota}_*$, respectively.

The above objective is optimized with respect to the boundary shape of the plasma, which is described in terms of Fourier modes. The boundary can be described by

$$\left. \begin{aligned} R(\theta, \phi) &= \sum_{m=0}^{M_{\text{pol}}} \sum_{n=-N_{\text{tor}}}^{N_{\text{tor}}} r_{c,m,n} \cos(m\theta - n_{\text{ip}}n\phi), \\ Z(\theta, \phi) &= \sum_{m=0}^{M_{\text{pol}}} \sum_{n=-N_{\text{tor}}}^{N_{\text{tor}}} z_{s,m,n} \sin(m\theta - n_{\text{ip}}n\phi), \end{aligned} \right\} \tag{2.4}$$

where R , Z , ϕ are the cylindrical coordinates of our surface, M_{pol} and N_{tor} are the maximum poloidal and toroidal mode numbers in our representation and n_{fp} is the number of field-periods of the geometry. Throughout this work, stellarator symmetry is assumed, which eliminates the sine and cosine components in R and Z , respectively. The poloidal angle θ is arbitrary. Typically, $M_{\text{pol}} = N_{\text{tor}} = 6$ in our optimizations.

For each boundary, a magnetic equilibrium is calculated using the VMEC code (Hirshman & Whitson 1983) and the objective is evaluated for this calculated equilibrium. The optimization is formulated using the SIMSOPT framework (Landreman *et al.* 2021). We use the Trust Region Reflective method in SCIPY to find a local minimum given an initial condition (Virtanen *et al.* 2020).

2.1. Computational procedure

We generate our equilibria by solving a series of local optimization problems with different \bar{l}_* , using a continuation method.

Let $F[\bar{l}_*]$ denote the objective function (2.2) with a specific value of \bar{l}_* . Let $F^{(0)} = F[\bar{l}_*^{(0)}]$ and let $\mathbf{x}^{(0)}$ be an optimum to this objective. We define our series of optimization problems through

$$\bar{l}_*^{(n+1)} = \bar{l}_*^{(n)} + \Delta\bar{l}_*, \quad (2.5)$$

$$F^{(n+1)} = F[\bar{l}_*^{(n+1)}]; \quad (2.6)$$

and take $\mathbf{x}^{(n)}$ to denote the optimum of $F^{(n)}$.

Since we are doing local optimizations, the obtained optimum for a given $F^{(n)}$ will depend on the initial guess. For each new optimization problem $F^{(n+1)}$, we use the optimum from the previous optimization problem $F^{(n)}$ as our initial guess. If $\Delta\bar{l}_*$ is small, we expect $F^{(n+1)}$ to be roughly equal to $F^{(n)}$ with a small perturbation, and $\mathbf{x}^{(n)}$ can thus be expected to be close to the new optimum $\mathbf{x}^{(n+1)}$. If this expectation holds, our procedure will generate a series of equilibria with different \bar{l}_* that are otherwise similar, in the sense of being close to each other in the space of Fourier boundary modes over which we optimize.

3. Results

In this section, we present measures of quasi-symmetry, fast-particle confinement and ion-temperature-gradient (ITG) stability for the different geometries obtained using the procedure described in § 2.1. Quasi-symmetry is part of the objective function (2.2), and can thus be considered a measure of how well the optimizer performs for the different \bar{l}_* . The other figures of merit were not explicitly optimized for.

3.1. Exploring the optimization landscape around an optimum

Before presenting detailed results of our optimized configurations, it is instructive to look at the optimization landscape around an optimum and how it changes when going to the next \bar{l}_* . In figure 1, we plot the objective function when varying each of the Fourier boundary modes about our optimum for $\bar{l}_* = 0.42$. Each panel corresponds to a different boundary mode, indicated above the plot. The y axes and x axes correspond to the value of the objective and the boundary mode, respectively. Colour is used to indicate the range of the y axis. The range of the x axis covers a 30% relative change in the corresponding boundary parameter.

From reading the colourbar, we see that the relative sensitivity of the objective to various boundary modes varies by almost eight orders of magnitude. This is largely due to the differences in the absolute values of the original modes. Since several boundary modes have very shallow or even non-existent local wells, our optimum is effectively in a flat

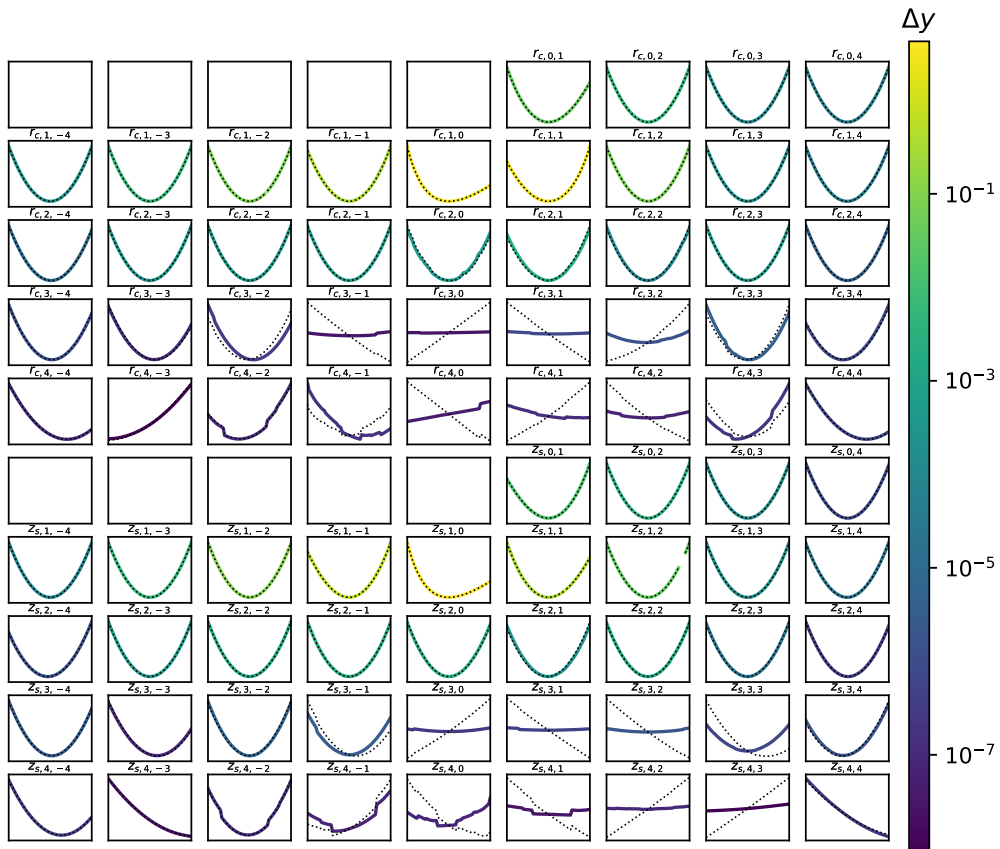


FIGURE 1. Solid lines: variation in the objective (2.2) for $\bar{l}_* = 0.42$ for a scan in boundary modes around the corresponding optimum. Each panel corresponds to a scan in one boundary mode, as indicated above the panel. Each boundary mode is varied by up to $\pm 30\%$ and the colours correspond to the range of values taken by the objective along that curve. Dashed lines: objective evaluated for the same boundaries, but with $\bar{l}_* = 0.43$.

valley and is thus not unique. Small changes in the numerical optimization method or initial condition will result in optima where the non-sensitive boundary modes assume different values, within the narrow range plotted.

This ‘non-uniqueness’ suggests that the slight perturbations of the continuation method may cause unpredictable jumps in the insensitive boundary modes. To illustrate the immediate effect of increasing \bar{l}_* , figure 1 also shows the landscape around the same $\bar{l}_* = 0.42$ optimum after \bar{l}_* in the objective has been changed to 0.43, corresponding to the initial state of the optimization of the next step in our \bar{l}_* scan. For some of the insensitive modes (for example, $m = 3, n = 0$), the objective changes from being relatively flat to having an incline over the plotted range. This will push the optimizer to make relatively large jumps in these boundary modes, potentially to different local minima in these insensitive modes.

Thus, the continuation method does not strictly follow a unique minimum as the landscape is deformed by changing \bar{l}_* . Rather, it samples a class of minima with potentially large relative differences in the insensitive modes. In the next section we use the resulting equilibria to evaluate how various figures of merit vary with \bar{l} .

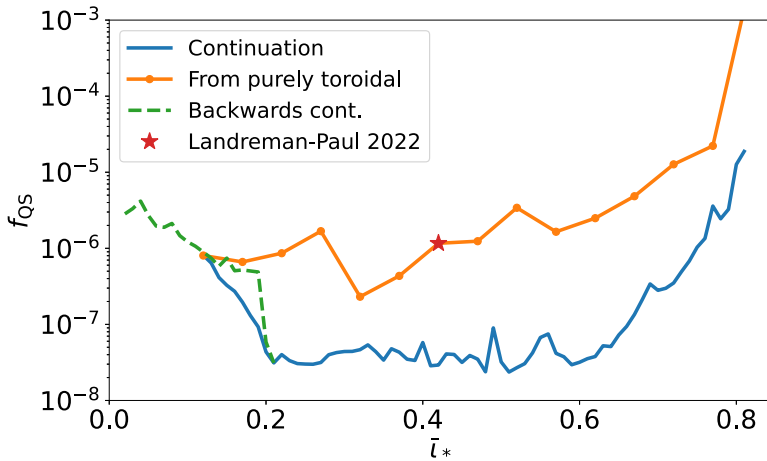


FIGURE 2. Quasi-symmetry error (2.3) from configurations with varying \bar{t} . The blue line corresponds to configurations found by the continuation method described in § 2.1. For comparison, we also include optimizations following the same procedure as in Landreman & Paul (2022), but targeting different \bar{t} (orange). Depicted in green is a continuation scan towards lower \bar{t} , showing that the scan is not reversible, which is expected given the ‘non-uniqueness’ of the configurations as discussed in § 3.1. The ‘Precise QA’ (Landreman & Paul 2022) corresponds to the $\bar{t} = 0.42$ ‘from purely toroidal’ optimization.

3.2. Varying rotational transform in a QA equilibrium

Using the objective function (2.2) with $A_* = 6.0$, $N = 1$ and $M = 0$ in (2.3), we optimize the boundary of an $n_{\text{fp}} = 2$ field-period vacuum configuration starting from a purely toroidal field and targeting $\bar{t}_* = 0.12$. We optimize by successively expanding the number of boundary Fourier modes in six steps, only including Fourier modes up to $|n| \leq j$, $m \leq j$ in the j th step. This optimization is repeated for six different relative finite-difference step sizes to calculate the rate of change with respect to the different boundary modes $r_{c,m,n}$ and $z_{s,m,n}$, and the optimum with the lowest quasi-symmetry error is picked. The resulting optimum is taken as our initial configuration in the procedure described in the next paragraph.

Incrementing \bar{t}_* by 0.01, we redo the optimization starting from the previous optimum. As this is a small adjustment to the objective function, the previous optimum is expected to be a good initial guess. Thus, we immediately optimize with all Fourier boundary modes. This not only saves time, but is also expected to make the resulting series of optima similar to each other. For convenience, all optimizations are performed with the initially found optimal finite-difference step size.

The quasi-symmetry errors of the resulting configurations at the different \bar{t}_* are plotted in figure 2, alongside the results of rerunning the optimizations from a purely toroidal field at different \bar{t}_* . We see that the continuation method results in lower quasi-symmetry error throughout the entire \bar{t}_* range, achieving quasi-symmetry errors about ten times lower than that for the optimization starting from a purely toroidal field. For $\bar{t}_* \in [0.2, 0.65]$, the quasi-symmetry error displays a broad valley, where it varies between 2.4×10^{-8} and 9.0×10^{-8} .

For $0.12 < \bar{t}_* < 0.2$, the continuation method gives decreasing quasi-symmetry error for increasing \bar{t}_* . To test whether this is due to lingering effects of the initial configuration

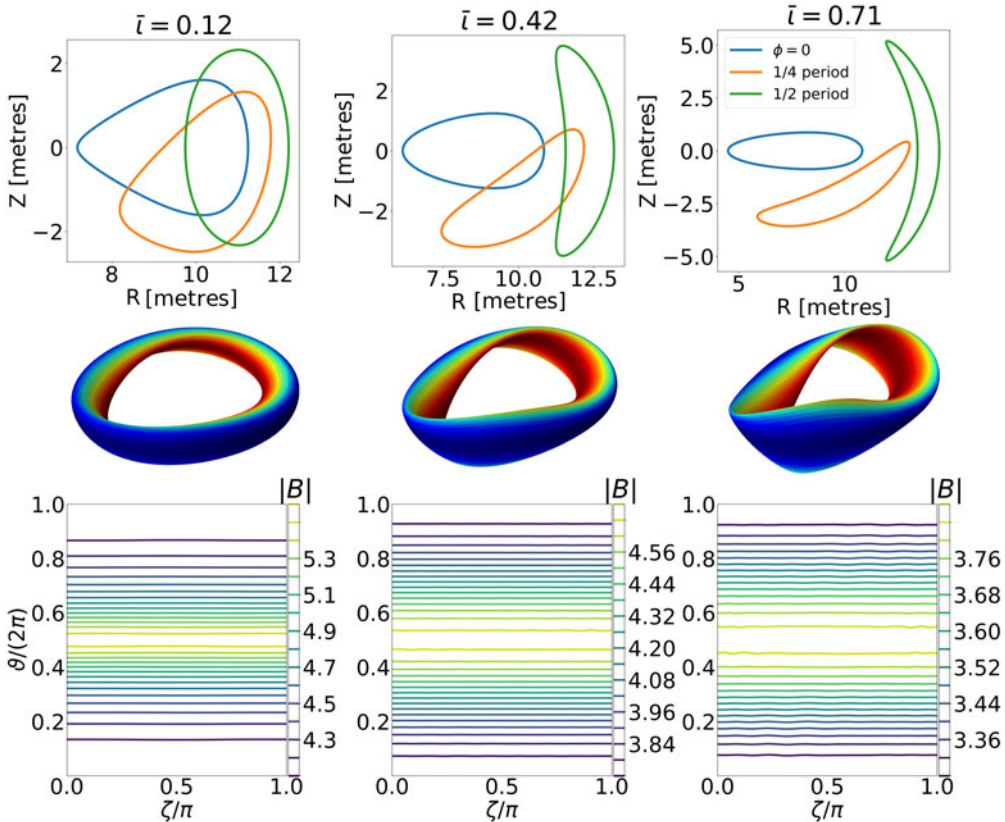


FIGURE 3. Rows 1 and 2: plasma boundary shapes for select points in the scan ($\bar{\iota} = 0.12$, $\bar{\iota} = 0.42$ and $\bar{\iota} = 0.71$). As $\bar{\iota}$ increases, the boundary becomes more elongated. Row 3: $|B|$ on the boundary in Boozer coordinates.

at $\bar{\iota}_* = 0.12$, we run the continuation method in reverse starting from the $\bar{\iota}_* = 0.21$ configuration. The result is shown in green in figure 2. Due to the ‘non-uniqueness’ property discussed in § 3.1, we do not expect the process to retrace itself exactly when done in reverse. Indeed, we find that the resulting configurations have somewhat different quasi-symmetry error after only two backwards iterations, compared with the previously obtained configurations with the same $\bar{\iota}_*$. These new configurations are not notably better than the configurations obtained by starting from a purely toroidal field. This indicates that higher quasi-symmetry error for $\bar{\iota}_* < 0.20$ is not merely an effect of the initial conditions in the optimization.

Figure 3 shows examples of boundary shapes obtained in the continuation scan. We see that as $\bar{\iota}_*$ increases, the boundary shapes become more elongated. This appears to limit the range of achievable $\bar{\iota}_*$ in the scan, as VMEC eventually fails to solve for the resulting plasma equilibria around $\bar{\iota}_* = 0.81$. The quasi-symmetry starts to degrade already at $\bar{\iota}_* \approx 0.65$, perhaps because of the optimum shape becoming thin enough that it cannot be accurately represented using only boundary Fourier modes up to $|n| \leq 4$, $m \leq 4$.

One advantage of higher $\bar{\iota}_*$ is that the orbit width of charged particles roughly scales as $1/\iota$ for QA configurations. Specifically, in the small-orbit-width limit, the orbit width is




FIGURE 4. Fast-particle loss fraction for the different \bar{l}_* configurations as calculated by SIMPLE by tracing 5000 alpha particles released at the surfaces with normalized toroidal flux $s = 0.3$, $s = 0.5$, $s = 0.7$. The dashed lines show least-squares fits of $c\bar{l}_*^\alpha$ to the data, with fit parameters c and α shown in the figure.

proportional to (Paul *et al.* 2022)

$$f_{\text{orbit width}} = \left| \frac{MG - NI}{Mt - N} \right|. \quad (3.1)$$

We thus expect fast particles to be better confined in QA configurations with higher \bar{l}_* , all else being equal.

This prediction is borne out by particle tracing simulations using the collisionless guiding-centre orbit solver SIMPLE (Albert, Kasilov & Kernbichler 2020*a,b*). Figure 4 shows the loss fractions calculated by tracing 5000 alpha particles with 3.52 MeV energy for 0.2 s. The particles were launched at three different radii, and considered lost when crossing the last closed flux surface at normalized toroidal flux $s = 1.0$. The configurations have all been scaled up to a reactor-relevant minor radius of 1.704 m and volume-average magnetic field of 5.865 T, to match the ARIES-CS configuration (Najmabadi *et al.* 2008). We see that the confinement improves with \bar{l}_* until about $\bar{l}_* = 0.73$, where the increase in quasi-symmetry error outweighs the benefits of increasing \bar{l}_* .

To see to which extent the $1/\iota$ orbit-width scaling is reflected in the fast-particle losses, we fit the losses to $c\bar{l}_*^\alpha$ for $\bar{l}_* \in [0.2, 0.65]$, where the quasi-symmetry error is relatively constant. The fits and values of the coefficients are shown alongside the data in figure 4, for the $s = 0.5$ and $s = 0.7$ results. The fits systematically overpredict the particle losses towards the higher end of the fitted range. Increasing the range of \bar{l}_* in the fit to $[0.12, 0.73]$ changes the values of the coefficients to $(c = 0.07, \alpha = -0.99)$ and $(c = 0.017, \alpha = -1.6)$ for the $s = 0.7$ and $s = 0.5$ fits, respectively, and does not significantly affect the agreement between the fitted curve and the data. Thus, the orbit width scaling with ι only offers a qualitative estimate of how the fast-particle losses scale, and other mechanisms must be accounted for to get the full picture.

Not included in the SIMPLE calculations is the effect of the gyro-orbit about the guiding centres, which effectively widens the orbits and thus should increase the losses. As an estimate of the size of this effect, we calculate the ratio of the gyroradius to the orbit width, using the expression for the orbit width in the small-orbit-width limit given by Paul




FIGURE 5. (a) Maximum growth rate at $s = 0.25$, for a different number N_α of flux tubes in each $\bar{\iota}$ configuration. Data are from electrostatic linear STELLA simulations with $a/L_T = 3$, $a/L_n = 1$ and adiabatic electrons. (b) The k_y mode number of the fastest growing mode. The underlying $\gamma(k_y)$ curves are shown in figure 7. Vertical dotted lines indicate some low-order rationals.

et al. (2022):

$$\frac{\text{gyroradius}}{\text{orbit width}} = \frac{|N - \iota M|}{|MG - NI|} |\nabla\psi| \frac{\sqrt{B\lambda}}{2\sqrt{1 - B\lambda}}, \quad (3.2)$$

where $\lambda = v_\perp^2 / (Bv^2)$, with v_\perp the speed of the particle in the direction perpendicular to the magnetic field. For our QA configurations ($I \approx 0$ Tm, $G \approx 70$ Tm, $-\nabla\psi \approx 20$ Tm, $N = 0$, $M = 1$) this ratio ends up being roughly $\iota/7$ for $B\lambda = 0.5$, and is thus much less than one. The finite gyro-orbit correction should thus be small for our range of ι , except for the most deeply trapped particles. Specifically, the gyro-orbit width becomes larger than the orbit width for $B\lambda > (7/\iota)^2 / [1 + (7/\iota)^2]$, or about $B\lambda > 0.987$ for the worst-case scenario of $\iota = 0.81$.

For configurations with this low quasi-symmetry error, collisional transport is almost surely dwarfed by turbulent transport. For example, ITG mode turbulence appears to dominate the heat transport in Wendelstein 7-X (Carralero et al. 2021).

In figure 5, we show the maximum growth rates $\max(\gamma)$ for the $s = 0.25$ surface in each configuration, as calculated from linear electrostatic gyrokinetic flux-tube simulations using the gyrokinetic code STELLA (Barnes, Parra & Landreman 2019) using adiabatic




FIGURE 6. Shear $\hat{s} = -(s/\iota)(d\iota/ds)$ at $s = 0.25$ plotted against the $\bar{\iota}$ of each equilibrium. Zero shear is marked with a dashed line.

electrons. The simulations use gradients to promote ITG turbulence: $a/L_{Ti} = 3$ and $a/L_n = 1$. In the figure, N_α indicates the number of flux tubes in which $\max(\gamma)$ was calculated, using equally spaced values of the field line label α_0 on the interval $[0, \pi]$, with a point always placed on $\alpha_0 = 0$. In STELLA, α_0 (together with s) specifies the location of a flux tube, with α_0 being the value of $\alpha = \theta - \iota\phi$ at the middle of the flux tube. The α_0 interval $[0, \pi]$ covers the entire half-period of a stellarator symmetric configuration. The simulations use $k_x = 0$ and the parallel domain extends for 70 toroidal turns, corresponding to about 8.4 poloidal turns for the lowest ι values. Due to the low shear of these configurations, the modes can have structure far along the field line; 70 toroidal turns was found to be sufficient to resolve this with good margins towards the boundaries for all configurations in the scan. The shear \hat{s} in each configuration at the radius of the simulations ($s = 0.25$) is shown in figure 6, revealing that most of the scan has a slight positive shear, with the configurations with $\bar{\iota} < 0.2$ having somewhat larger shear still below 0.1. Note that while 0.1 is small in an absolute sense, it is not unusual for stellarators to have such low shear (Ascásbar *et al.* 2008). We only consider $k_x = 0$ modes, since these modes usually have the largest growth rates for ITG turbulence, which we confirmed for the $\bar{\iota} = 0.49$ and $\bar{\iota} = 0.49$ cases.

The $\gamma(k_y)$ curves corresponding to the flux tube with the maximum growth rates are shown in figure 7, for some of the $\bar{\iota}_*$ configurations. Despite some of the curves being calculated on flux tubes with different α_0 , the result is an essentially smooth change in the growth rate curves with $\bar{\iota}_*$. For $\bar{\iota}_* \geq 0.77$, the growth rate peaks above $\rho k_y = 3.0$, so these simulations were modified to include higher k_y values.

Figures 5 and 7 both show a clear trend towards higher $\max(\gamma)$ and higher wavenumbers k_y with increasing $\bar{\iota}$, suggesting smaller-scale turbulence. At certain values of $\bar{\iota}$, the $N_{\alpha 0} = 1$ and $N_{\alpha 0} = 2$ curves in figure 5 display sharp dips towards lower $\max(\gamma)$.

This can be understood in terms of the flux tubes on rational surfaces not sampling the full flux-surface geometry. As a result, modes on different flux tubes will, for example, experience different levels of finite-Larmor-radius (FLR) damping in regions of bad curvature. This is illustrated in figure 8, which shows the value of $|\nabla\alpha|^2$ at maxima in the curvature drift along the flux tube, for different flux tubes in the $\bar{\iota}_* = 0.49$ and $\bar{\iota}_* = 0.50$ geometries. For the non-rational ι , all flux tubes sample regions of bad curvature with all values of $|\nabla\alpha|^2$, while for the rational ι , the different flux tubes display different values of $|\nabla\alpha|^2$ at regions of bad curvature. Flux tubes with lower $|\nabla\alpha|^2$ in bad-curvature regions have higher growth rates as they experience less FLR damping.

Figure 9 shows the corresponding nonlinear ion heat flux as calculated using the gyrokinetic code GX. GX is a new gyrokinetic code that runs on GPUs (Mandell, Dorland & Landreman 2018; Mandell *et al.* 2022). GX has been benchmarked against STELLA




FIGURE 7. Growth rates $\gamma(k_y)$ for the flux tube with the highest max(γ) for a subset of the different $\bar{\iota}_*$ configurations. The curves for different configurations are grouped into different panels based on the range of their $\bar{\iota}_*$ values: (a) 0.13–0.34, (b) 0.37–0.56 and (c) 0.59–0.77. Colours indicate the specific $\bar{\iota}_*$ corresponding to each curve.

and other established gyrokinetic codes (Mandell *et al.* 2022), and should be faster than STELLA when it comes to running electrostatic nonlinear gyrokinetic simulations with adiabatic electrons. In contrast to the linear results, there is more variation between flux tubes with different α_0 , but seemingly no great difference between rational and irrational values of ι . Both these features are consistent with the nonlinear heat flux being insensitive to structures that are very elongated along the field line, so that even on irrational flux tubes, the turbulence does not effectively sample all regions of bad curvature on the flux surface. The trend with respect to ι is non-monotonic with several local maxima and minima.

We also investigated how far away coils can be placed from our plasma, using the method in Kappel *et al.* (2024). Given a maximum surface-current density K_{\max} (corresponding to a minimum coil–coil distance) and a required accuracy in reproducing the last-closed flux surface of a target configuration, the distance between the plasma and the coil winding surface in the coil-shape code REGCOIL (Landreman 2017) is increased until the maximum current density on the winding surface reaches K_{\max} . The accuracy with which the coils produce the target last-closed flux surface is measured by

$$B_{n,\text{RMS}} = \sqrt{\frac{\int_{\text{LCFS}} dS (\mathbf{n} \cdot \mathbf{B})^2}{\int_{\text{LCFS}} dS}}, \tag{3.3}$$

where the surface integrals are over the target surface and \mathbf{n} is the unit normal of this surface; \mathbf{B} is here the magnetic field produced by the coils.

The results of REGCOIL calculations with $K_{\max} = 17.16 \text{ MA m}^{-1}$ and $B_{n,\text{RMS}} = 0.01 \text{ T}$ for the different $\bar{\iota}_*$ target configurations are shown in figure 10. All target configurations have been rescaled to the ARIES-CS minor radius a and volume-averaged B , and the value of K_{\max} was chosen to match the minimum coil–coil distance of the ARIES-CS coil set.

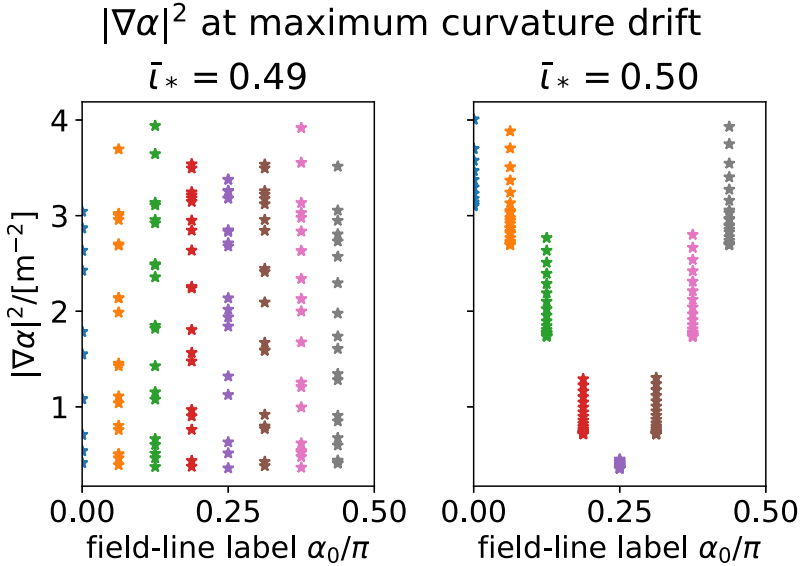


FIGURE 8. Values of $|\nabla\alpha|^2$ at points of maximum curvature drift along different flux tubes (given by α_0 on the x axis) for an irrational ($\bar{l}_* = 0.49$) and rational ($\bar{l}_* = 0.50$) flux surface.




FIGURE 9. Nonlinear ion heat fluxes Q_i for a few different \bar{l}_* configurations starting from $\bar{l}_* = 0.20$ to $\bar{l}_* = 0.80$. The fluxes are calculated using the gyrokinetic code GX on a single flux tube, with calculations done for the $\alpha_0 = 0$ and $\alpha_0 = \pi/4$ flux tubes. Here $Q_{gB} = n_i T_i v_{Ti}^3 / (B^2 a^2)$, where $v_{Ti} = \sqrt{2T_i/m_i}$ is defined with a $\sqrt{2}$ unlike in the raw GX output.

Also included in the figure is the gradient scale length of B , defined as

$$L_{\nabla B}^* \equiv \min \left(\frac{\sqrt{2}B}{\|\nabla\mathbf{B}\|_F} \right), \tag{3.4}$$

where $\|\nabla\mathbf{B}\|_F$ is the Frobenius norm of the tensor $\nabla\mathbf{B}$ and the minimum is taken over boundary surface of the plasma. This metric, which is purely defined in terms of the target magnetic field and requires no coil geometry to calculate, has been proposed as a good proxy for the plasma–coil separation. Again, see Kappel *et al.* (2024) for details.




FIGURE 10. Coil–plasma separation for the different $\bar{\iota}_*$ configurations, calculated using REGCOIL, alongside the magnetic scale length $L_{\parallel \nabla B}$. The configurations have been scaled up to $a = 1.704\text{m}$ and volume-average magnetic field of 5.865 T , to match ARIES-CS. The sheet current density is computed to meet a target accuracy $B_{n,\text{RMS}} = 0.01\text{ T}$ and a maximum current density of 17.16 MA m^{-1} , which corresponds to the minimum coil–coil distance of ARIES-CS (Kappel, Landreman & Malhotra 2024).

From figure 10, we see that the coil–plasma separation and gradient scale length both vary weakly over the ι scan. Specifically, the coil–plasma separation is between 2.85 and 3.25 m with a minimum for $\bar{\iota}_*$ around 0.35 . The difference between the smallest plasma–coil separation and the largest is only 23% of the minor radius. This is a small difference compared with the full range of plasma–coil separations present in Kappel *et al.* (2024). The $L_{\nabla B}^*$ metric is not expected to be a detailed enough proxy to fully capture such slight variations, which is reflected in the fact that $L_{\nabla B}^*$ does not reproduce the minimum at $\bar{\iota}_* = 0.35$.

We thus conclude that changes to $\bar{\iota}$ do not seem to play a large role in setting the plasma–coil separation. Since our configurations differ widely in their elongation, it also seems like elongation does not have a large effect on plasma–coil separation.

3.3. Varying shear in QA equilibrium

All configurations in the previous section have low values of global magnetic shear. This might make them susceptible to curvature-driven instabilities, compared with configurations with negative shear (Antonsen *et al.* 1996; Nadeem, Rafiq & Persson 2001).

High shear helps limiting the parallel extent of modes elongated along the magnetic field lines due to radially nearby field lines separating as the parallel coordinate is traversed, which forces such modes to have lower radial correlation length or a limited parallel extent. Gyrokinetic modes with smaller perpendicular extent generally have a smaller impact on a configuration, as they lead to less transport, which can be qualitatively understood by viewing the size of the mode perpendicular to the field line as a step length in a diffusive process (Merz & Jenko 2008).

We obtain configurations with stronger shear by additional optimizations where we add a term targeting the mean shear. Defining the mean shear as

$$\bar{s} = -c_1/\bar{\iota}, \quad (3.5)$$

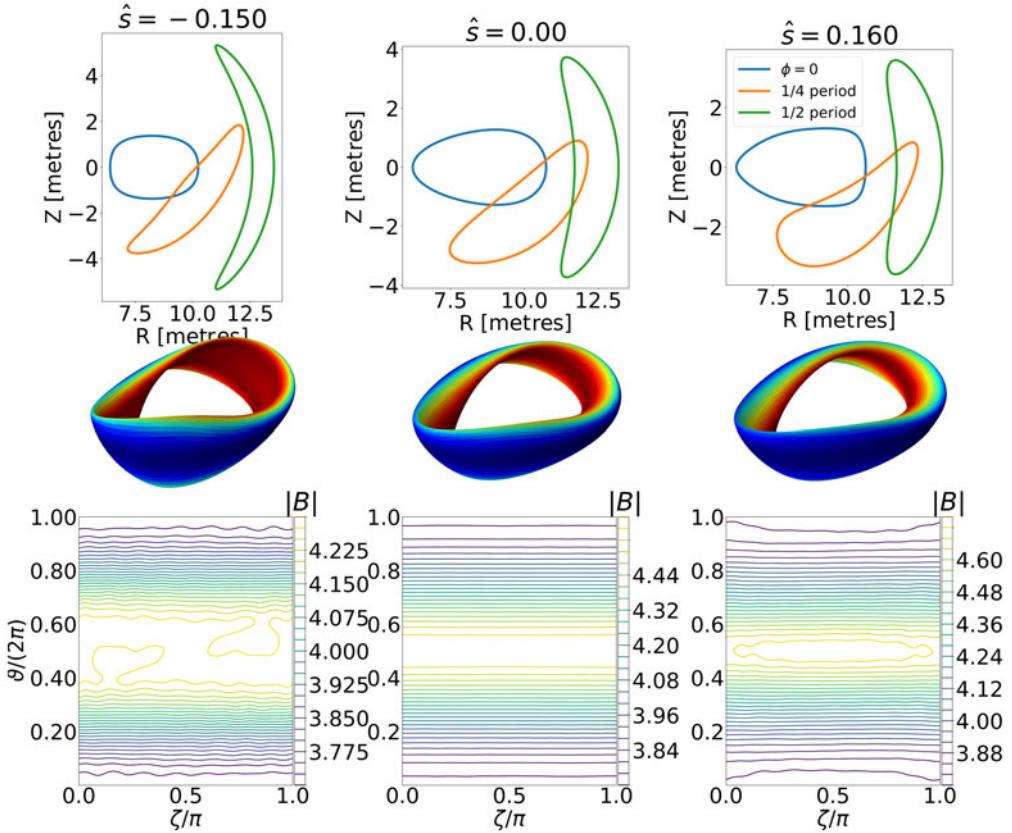


FIGURE 11. Plasma boundary shapes for select points in the scan ($\bar{s} = -0.16$, $\bar{s} = 0.0$ and $\bar{s} = 0.15$). Imposing positive shear mainly affects the triangularity of the boundary, while a negative shear makes the boundary more elongated and narrow.

where c_1 is the result of a linear fit $c_0 + c_1s$ to the $\iota(s)$ profile, we add the following optimization target to our objective (2.2):

$$(\bar{s} - \bar{s}_*)^2. \tag{3.6}$$

The negative sign in (3.5) is included to match the sign convention of the conventional definition of shear, $-(s/\iota)(d\iota/ds)$.

Starting from the $\bar{\iota}_* = 0.42$ configuration found in § 3.2, which has a mean shear $\bar{s} = 0.0195 \approx 0.02$, we do two continuation scans in \bar{s}_* towards more negative and more positive shear, respectively, still targeting $\bar{\iota}_* = 0.42$. The configurations with lowest and highest shear are shown in figure 11, alongside the $\bar{s} = 0$ configuration.

The resulting quasi-symmetry error is shown in figure 12. We see that imposing an additional shear objective degrades the quasi-symmetry. This is expected when adding competing objectives to an optimization problem.

In figure 13, we show the fast-particle loss fraction calculated using the same set-up as for the $s = 0.3$ calculations in figure 4. The degradation in quasi-symmetry appears to increase the fast-particle losses somewhat, but the effect is small enough for this range of shear for the uncertainty in the individual calculations to be comparable, as seen by the non-smoothness of the curve.




FIGURE 12. Quasi-symmetry error (2.3) from a continuation scan varying shear. The scan is performed around the $\iota = 0.42$ configuration.




FIGURE 13. Fast-particle loss fraction for the different \hat{s} configurations as calculated by SIMPLE using the same simulation set-up as for the $s = 0.3$ curve in figure 4.

Finally, we investigate the effect of shear on turbulent transport. The maximum growth rates $\max(\gamma)$ and the dependence of the growth rate on k_y for the shear scan are shown in figures 14 and 15, respectively. We see a clear trend towards increasing maximum growth rates for shear above $\hat{s} > 0.085$. This increase seems to be due to some new mode appearing at high k_y values, and the maximum growth rate of this mode is not properly resolved in these simulations. The mode remains even when setting the density gradient to zero, and must thus be driven by the ITG, although it appears at much larger values of k_y than typical ITG turbulence.

The effect on the nonlinear ion heat flux is more modest, but there is an increase of about 25% for positive shear, as seen in figure 16. Positive shear being destabilizing is well known in the tokamak literature (Antonsen *et al.* 1996; Nadeem *et al.* 2001), and our results appear to be consistent with that.




FIGURE 14. Maximum growth rates for the different \hat{s} configurations, using the same simulation set-up as in figure 5.




FIGURE 15. Growth rates $\gamma(k_y)$ for a subset of the different \hat{s}_* configurations.




FIGURE 16. Nonlinear heat fluxes for the different \bar{s} configurations.

4. Conclusions

We have used a continuation method to generate a sequence of QA equilibria with different rotational transform profiles. The configurations are similar to each other, with each a local minimum of an optimization problem that is initialized at the previously found minimum. Through this procedure, we have found $n_{fp} = 2$ QA configurations with higher mean rotational transform $\bar{\iota}$ than what was previously found. We find QA configurations with exceptionally low quasi-symmetry errors for $\bar{\iota} \in [0.2, 0.65]$. For $\bar{\iota}$ in and slightly above this range, increasing $\bar{\iota}$ increases the elongation of the plasma boundary, which is likely what ultimately limits the achievable values of $\bar{\iota}$. Fast-particle losses as calculated by a collisionless drift-kinetic particle-tracing code decrease with ι up to a point, but the effect is weaker than the ι^{-1} scaling predicted by purely accounting for how the orbit width scales with ι . Nevertheless, there is a potential conflict between having low elongation and low fast-particle losses. The dependence of the turbulent ion heat flux is non-monotonic in ι and shows some mild sensitivity to which flux tube the simulations are performed at. The linearly calculated growth rates, on the other hand, are sensitive to the choice of flux tube on rational flux surfaces.

Higher shear can be imposed on the configurations by adding an extra term in the objective function. This increases the quasi-symmetry error somewhat, but this hardly affects the fast-particle losses. The turbulent heat flux tends to increase somewhat for positive shear, causing an increase by about 25% for the range of shear considered here. The effect of shear in magnetohydrodynamic stability is likely even more important, but was not considered here.

Acknowledgements

The authors are grateful for discussions with B. Dorland, I. Abel, R. Jorge, E. Rodriguez, W. Sengupta, N. Nikulsin, R. Nies, T. Adkins and J. Parisi, who have all provided valuable input and motivation for this work.

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

Funding

This work was supported by the US Department of Energy, Office of Science, Office of Fusion Energy Science, under award number DE-FG02-93ER54197. This research used

resources of the National Energy Research Scientific Computing Center (NERSC), a US Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory, operated under Contract No. DE-AC02-05CH11231 using NERSC award FES-ERCAP-mp217-2023. Additional computations were performed on the HPC systems Cobra and Raven at the Max Planck Computing and Data Facility (MPCDF).

Declaration of interest

The authors report no conflict of interest.

REFERENCES

- ALBERT, C.G., KASILOV, S.V. & KERNBICHLER, W. 2020a Accelerated methods for direct computation of fusion alpha particle losses within, stellarator optimization. *J. Plasma Phys.* **86** (2), 815860201. <https://doi.org/10.1017/S0022377820000203>
- ALBERT, C.G., KASILOV, S.V. & KERNBICHLER, W. 2020b Symplectic integration with non-canonical quadrature for guiding-center orbits in magnetic confinement devices. *J. Comput. Phys.* **403**, 109065. <https://www.sciencedirect.com/science/article/abs/pii/S0021999119307703?via%3Dihub>
- ALLGOWER, E.L. & GEORG, K. 2003 *Introduction to Numerical Continuation Methods*. Society for Industrial and Applied Mathematics. <https://epubs.siam.org/doi/book/10.1137/1.9780898719154>
- ANTONSEN, T.M. JR., DRAKE, J.F., GUZDAR, P.N., HASSAM, A.B., LAU, Y.T., LIU, C.S. & NOVAKOVSKII, S.V. 1996 Physical mechanism of enhanced stability from negative shear in tokamaks: implications for edge transport and the L-H transition. *Phys. Plasmas* **3** (6), 2221–2223. <https://doi.org/10.1063/1.871928>
- ASCASÍBAR, E., *et al.* 2008 Effect of rotational transform and magnetic shear on confinement of stellarators. *Plasma Fusion Res.* **3**, S1004–S1004.
- BARNES, M., PARRA, F.I. & LANDREMAN, M. 2019 stella: an operator-split, implicit–explicit δf -gyrokinetic code for general magnetic field configurations. *J. Comput. Phys.* **391**, 365–380. <https://www.sciencedirect.com/science/article/abs/pii/S002199911930066X?via%3Dihub>
- BINDEL, D., LANDREMAN, M. & PADIDAR, M. 2023 Understanding trade-offs in stellarator design with multi-objective optimization. <https://arxiv.org/abs/2304.08698>
- BOOZER, A.H. 1983 Transport and isomorphic equilibria. *Phys. Fluids* **26** (2), 496–499. <https://doi.org/10.1063/1.864166>
- BULLER, S. 2024 Dataset for ‘A family of quasi-axisymmetric stellarators with varied rotational transform data’. <https://zenodo.org/records/10521394>
- CARRALERO, D., ESTRADA, T., MARAGKOUKAKIS, E., WINDISCH, T., ALONSO, J.A., BEURSKENS, M., BOZHENKOV, S., CALVO, I., DAMM, H., FORD, O., FUCHERT, G., GARCÍA-REGAÑA, J.M., PABLANT, N., SÁNCHEZ, E., PASCH, E., VELASCO, J.L. & THE WENDELSTEIN 7-X TEAM 2021 An experimental characterization of core turbulence regimes in Wendelstein 7-X. *Nucl. Fusion* **61** (9), 096015. <https://doi.org/10.1088/1741-4326/ac112f>
- CONLIN, R., DUDT, D.W., PANICI, D. & KOLEMEN, E. 2023 The DESC stellarator code suite. Part 2. Perturbation and continuation methods. *J. Plasma Phys.* **89** (3), 955890305. <https://doi.org/10.1017/S0022377823000399>
- GARREN, D.A. & BOOZER, A.H. 1991 Existence of quasi-helically symmetric stellarators. *Phys. Fluids B: Plasma Phys.* **3** (10), 2822–2834. <https://doi.org/10.1063/1.859916>
- 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** (4), 042511. <https://doi.org/10.1063/5.0129716>
- HELANDER, P. & SIMAKOV, A.N. 2008 Intrinsic ambipolarity and rotation in stellarators. *Phys. Rev. Lett.* **101**, 145003. <https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.101.145003>
- HIRSHMAN, S.P. & WHITSON, J.C. 1983 Steepest–descent moment method for three-dimensional magnetohydrodynamic equilibria. *Phys. Fluids* **26** (12), 3553–3568. <https://pubs.aip.org/aip/pfl/article-abstract/26/12/3553/809269/Steepest-descent-moment-method-for-three?redirectedFrom=fulltext>

- KAPPEL, J., LANDREMAN, M. & MALHOTRA, D. 2024 The magnetic gradient scale length explains why certain plasmas require close external magnetic coils. *Plasma Phys. Control. Fusion* **66** (2), 025018. <https://doi.org/10.1088/1361-6587/ad1a3e>
- LANDREMAN, M. 2017 An improved current potential method for fast computation of stellarator coil shapes. *Nucl. Fusion* **57** (4), 046003. <https://doi.org/10.1088/1741-4326/aa57d4>
- LANDREMAN, M. 2022 Mapping the space of quasisymmetric stellarators using optimized near-axis expansion. *J. Plasma Phys.* **88** (6), 905880616. <https://doi.org/10.1017/S0022377822001258>
- 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** (65), 3525. <https://doi.org/10.21105/joss.03525>
- LANDREMAN, M. & PAUL, E. 2022 Magnetic fields with precise quasisymmetry for plasma confinement. *Phys. Rev. Lett.* **128**, 035001. <https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.128.035001>
- LANDREMAN, M., SENGUPTA, W. & PLUNK, G.G. 2019 Direct construction of optimized stellarator shapes. Part 2. Numerical quasisymmetric solutions. *J. Plasma Phys.* **85** (1), 905850103. <https://doi.org/10.1017/S0022377818001344>
- MANDELL, N.R., DORLAND, W., ABEL, I., GAUR, R., KIM, P., MARTIN, M. & QIAN, T. 2022 GX: a GPU-native gyrokinetic turbulence code for tokamak and stellarator design. <https://arxiv.org/abs/2209.06731>
- MANDELL, N.R., DORLAND, W. & LANDREMAN, M. 2018 Laguerre–Hermite pseudo-spectral velocity formulation of gyrokinetics. *J. Plasma Phys.* **84** (1), 905840108. <https://doi.org/10.1017/S0022377818000041>
- MERZ, F. & JENKO, F. 2008 Nonlinear saturation of trapped electron modes via perpendicular particle diffusion. *Phys. Rev. Lett.* **100**, 035005. <https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.100.035005>
- MOBAHI, H. & FISHER, J. III 2015 A theoretical analysis of optimization by Gaussian continuation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 29. Association for the Advancement of Artificial Intelligence. <https://doi.org/10.1609/aaai.v29i1.9356>
- NADEEM, M., RAFIQ, T. & PERSSON, M. 2001 Local magnetic shear and drift waves in stellarators. *Phys. Plasmas* **8** (10), 4375–4385. <https://doi.org/10.1063/1.1396842>
- NAJMABADI, F., *et al.* 2008 The ARIES-CS compact stellarator fusion power plant. *Fusion Sci. Technol.* **54** (3), 655–672. <https://doi.org/10.13182/FST54-655>
- NÜHRENBERG, J. & ZILLE, R. 1988 Quasi-helically symmetric toroidal stellarators. *Phys. Lett. A* **129** (2), 113–117. <https://www.sciencedirect.com/science/article/abs/pii/0375960188900801?via%3Dihub>
- PAUL, E.J., BHATTACHARJEE, A., LANDREMAN, M., ALEX, D., VELASCO, J.L. & NIES, R. 2022 Energetic particle loss mechanisms in reactor-scale equilibria close to quasisymmetry. *Nucl. Fusion* **62** (12), 126054. <https://doi.org/10.1088/1741-4326/ac9b07>
- PLUNK, G.G. & HELANDER, P. 2018 Quasi-axisymmetric magnetic fields: weakly non-axisymmetric case in a vacuum. *J. Plasma Phys.* **84** (2), 905840205. <https://doi.org/10.1017/S0022377818000259>
- PLUNK, G.G., LANDREMAN, M. & HELANDER, P. 2019 Direct construction of optimized stellarator shapes. Part 3. Omnigenity near the magnetic axis. *J. Plasma Phys.* **85** (6), 905850602. <https://doi.org/10.1017/S002237781900062X>
- RODRIGUEZ, E., HELANDER, P. & BHATTACHARJEE, A. 2020 Necessary and sufficient conditions for quasisymmetry. *Phys. Plasmas* **27** (6), 062501. <https://doi.org/10.1063/5.0008551>
- VIRTANEN, P., *et al.* 2020 SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nat. Meth.* **17**, 261–272. <https://doi.org/10.1038/s41592-019-0686-2>