Direct computation of magnetic surfaces in Boozer coordinates and coil optimization for quasi-symmetry

Andrew Giuliani[1]†, Florian Wechsung[1], Georg Stadler[1], Antoine Cerfon[1], and Matt Landreman[2]

[1]Courant Institute of Mathematical Sciences, New York University, New York, NY, USA
[2]University of Maryland-College Park, MD, USA

(Received xx; revised xx; accepted xx)

We propose a new method to compute magnetic surfaces that are parametrized in Boozer coordinates for vacuum magnetic fields. We also propose a measure for quasi-symmetry on the computed surfaces and use it to design coils that generate a magnetic field that is quasi-symmetric on those surfaces. The rotational transform of the field and complexity measures for the coils are also controlled in the design problem. Using an adjoint approach, we are able to obtain analytic derivatives for this optimization problem, yielding an efficient gradient-based algorithm. Starting from an initial coil set that presents nested magnetic surfaces for a large fraction of the volume, our method converges rapidly to coil systems generating fields with excellent quasi-symmetry and low particle losses. In particular for low complexity coils, we are able to significantly improve the performance compared to coils obtained from the standard two-stage approach, e.g. reduce losses of fusion-produced alpha particles born at half-radius from 17.7% to 6.6%. We also demonstrate 16-coil configurations with alpha loss $< 1\%$ and neoclassical transport magnitude $\epsilon_{\text{eff}}^{3/2}$ less than approximately 5×10^{-9}.

1. Introduction

Single-stage coil design, in which one optimizes for the geometry and currents of electromagnetic coils at the same time as the target magnetic configuration, is a promising approach for stellarator design as it considers physics goals and engineering constraints simultaneously (Giuliani *et al.* 2020; Henneberg *et al.* 2021; Wechsung *et al.* 2021; Yu *et al.* 2022). It has the potential to yield designs with coil systems that are easier to manufacture and position. In the single-stage method, the optimization objective must explicitly favor both the existence of nested flux surfaces for a large fraction of the plasma volume, as well as good particle confinement. This is empirically known to be more complicated than in the more common two-stage coil design strategy (Merkel 1987; Drevlak 1998; Strickler *et al.* 2002, 2004; Brown *et al.* 2015; Zhu *et al.* 2017; Landreman 2017; Paul *et al.* 2018; Singh *et al.* 2020). One single-stage optimization approach was demonstrated in Giuliani *et al.* (2020), based on the near-axis expansion for quasi-symmetry in Garren & Boozer (1991); Landreman & Sengupta (2018); Landreman *et al.* (2019); Landreman & Sengupta (2019). This approach was motivated by the fact that quasi-symmetric magnetic fields are known to have good confinement properties (Helander 2014; Landreman & Paul 2022). We demonstrated that this single-stage method effectively produced vacuum quasi-symmetric magnetic fields in the vicinity of the magnetic axis. However, we also observed

† Email address for correspondence: giuliani@cims.nyu.edu

that quasi-symmetry degraded with distance from the magnetic axis. The purpose of the present article is to introduce an optimization objective that targets quasi-symmetry at an arbitrary number of surfaces, potentially far from the magnetic axis. We show that this formulation results in coil systems producing a vacuum field with excellent quasi-symmetry for a large fraction of the plasma volume.

In order to optimize for quasi-symmetry on a given surface, we need to measure the degree of departure of the magnetic field from quasi-symmetry, which we do using Boozer coordinates (Boozer 1981). The first contribution of the present article thus is a new numerical method for computing magnetic surfaces directly parametrized in Boozer coordinates. It relies on a partial differential equation (PDE) directly derived from to contra- and covariant expressions for the field in Boozer coordinates. This PDE can be discretized either as a balanced or over-determined system of nonlinear equations, solved exactly or in a least squares sense, respectively.

The second contribution of this article is the formulation of a single-stage optimization problem for highly accurate quasi-symmetry and a target rotational transform on an arbitrary number of surfaces subject to engineering constraints for the coils. Penalty terms are constructed to favor simpler coils by constraining their minimum pairwise distance, length, maximum curvature, and mean squared curvature, similar to constraints used for two-stage coil design (Zhu *et al.* 2017; Kruger *et al.* 2021; Singh *et al.* 2020). We consider zero-thickness filament coils for simplicity, but our approach can be generalized to coils of finite thickness, e.g., following the method proposed in Singh *et al.* (2020).

We rely on gradient-based optimization tools for efficient and robust convergence to minimizers. To compute the required gradients in an efficient and accurate manner, we use an adjoint sensitivity approach. As a result, the performance of the method to compute gradients is independent of the number of coil degrees of freedom, which is of the order of hundreds in our example.

We apply our method to improve a set of coils obtained from the classic two-stage approach. These coils were optimized to approximate the precise quasi-symmetric magnetic fields recently discovered in Landreman & Paul (2022). Our approach allows us to improve quasi-symmetry on surfaces compared to the results from the two-stage approach, without increasing our measures for coil complexity, i.e., maximum curvature, mean squared curvature, minimum pairwise coil distance, and total coil length. For this particular example, our approach can be viewed either as a third stage improving an existing design found from the standard two-stage approach, or as single stage coil optimization approach using a good initialization for the coils.

The resulting configurations have impressive confinement, even though they only rely on 16 coils. When the three longest coil configurations are scaled to a reactor, losses of collisionless fusion-produced alpha particles born at half-radius are found to be below 1%. Moreover, the magnitude of the 3D neoclassical transport metric $\epsilon_{\text{eff}}^{3/2}$ (Nemov *et al.* 1999) is less than approximately 5×10^{-9}. To our knowledge this value is smaller than in any previous stellarator, smaller even than in tokamaks (Figure 15 of Spong (2015)).

The structure of this article is as follows. Sections 2 and 3 focus on our new numerical method for computing surfaces in Boozer coordinates: in Section 2, we introduce a parametrization to approximate toroidal surfaces, and in Section 3 we introduce the partial differential equations that the surfaces in Boozer coordinates satisfy, and the numerical method used to solve them. In Section 4, we define the coil optimization problem, which combines our surface computation approach, the physics goals of quasi-symmetry and a target rotational transform, with standard engineering constraints. In

Section 5, we study the performance of our new optimization scheme for the design of a quasi-axisymmetric magnetic configuration, and we summarize our work in Section 6.

2. Surface parametrization

We approximate a generic toroidal surface $\boldsymbol{\Sigma}$ that satisfies n_{fp}-rotational symmetry with a finite-dimensional surface $\boldsymbol{\Sigma}_s(\varphi,\theta) : [0,1)^2 \to (x,y,z)$. The parametrization of $\boldsymbol{\Sigma}_s(\varphi,\theta)$ detailed below directly incorporates the discrete rotational symmetry, making it efficient and convenient to work with. For $m_{\mathrm{pol}}, n_{\mathrm{tor}} \geqslant 0$, it is given by

$$\begin{aligned} x &= \cos(2\pi\varphi)\hat{x} - \sin(2\pi\varphi)\hat{y} \\ y &= \sin(2\pi\varphi)\hat{x} + \cos(2\pi\varphi)\hat{y} \\ z &= \sum_{i=0}^{2m_{\mathrm{pol}}} \sum_{j=0}^{2n_{\mathrm{tor}}} z_{i,j} w_i(\theta) v_j(\varphi) \end{aligned}$$

with

$$\hat{x} = \sum_{i=0}^{2m_{\mathrm{pol}}} \sum_{j=0}^{2n_{\mathrm{tor}}} x_{i,j} w_i(\theta) v_j(\varphi) \quad \text{and} \quad \hat{y} = \sum_{i=0}^{2m_{\mathrm{pol}}} \sum_{j=0}^{2n_{\mathrm{tor}}} y_{i,j} w_i(\theta) v_j(\varphi)$$

and the basis functions

$$\begin{aligned} w(\theta) &= (1, \cos(2\pi\theta), \sin(2\pi\theta), \ldots, \cos(2\pi m_{\mathrm{pol}}\theta), \sin(2\pi m_{\mathrm{pol}}\theta)) \\ v(\varphi) &= (1, \cos(2\pi n_{\mathrm{fp}}\varphi), \sin(2\pi n_{\mathrm{fp}}\varphi), \ldots, \cos(2\pi n_{\mathrm{fp}} n_{\mathrm{tor}}\theta), \sin(2\pi n_{\mathrm{fp}} n_{\mathrm{tor}}\theta)). \end{aligned}$$

The functions $\hat{x}$, $\hat{y}$ and z are linear combinations of tensor-product basis functions, which are n_{fp}-periodic in φ. In order to achieve rotational symmetry around the $\{x = y = 0\}$ axis, we apply a rotation to $\hat{x}$ and $\hat{y}$, which ensures that $\boldsymbol{\Sigma}_s$ is n_{fp}-rotational symmetric, as we show explicitly in Appendix A.

For stellarator-symmetric surfaces, we additionally require that

$$\begin{aligned} x(\varphi,\theta) &= x(-\varphi,-\theta) \\ y(\varphi,\theta) &= -y(-\varphi,-\theta) \\ z(\varphi,\theta) &= -z(-\varphi,-\theta). \end{aligned} \tag{2.1}$$

This is achieved by removing the basis functions in the x, y, z expansions that do not satisfy (2.1). In particular, basis functions of the form $\sin(2m\pi n_{\mathrm{fp}}\varphi)\cos(2n\pi\theta)$ and $\cos(2m\pi n_{\mathrm{fp}}\varphi)\sin(2n\pi\theta)$ are removed from $\hat{x}$. Similarly, basis functions of the form 1, $\sin(2m\pi n_{\mathrm{fp}}\varphi)\sin(2n\pi\theta)$, and $\cos(2mn_{\mathrm{fp}}\pi\varphi)\cos(2n\pi\theta)$ are removed from $\hat{y}$ and z. This results in a surface representation that has $n_s = 3[2n_{\mathrm{tor}}m_{\mathrm{pol}} + m_{\mathrm{pol}} + n_{\mathrm{tor}} + 1] - 2$ degrees of freedom. In the remainder of this manuscript we only consider stellarator-symmetric surfaces. Note that the work described below is not fundamentally restricted to this surface parametrization and others can be used. For example, we have also used a surface parametrization where the geometric degrees of freedom correspond to pointwise solution values and spatial derivatives are obtained using spectral differentiation matrices.

Note that the surface parametrization used, e.g., in VMEC (Hirshman & Whitson 1983) uses a parameterization in cylindrical coordinates $\boldsymbol{\Sigma}_s(\phi,\theta) = (r(\phi,\theta), \phi, z(\phi,\theta))$. This parameterization has fewer degrees of freedom to represent surfaces, but it is inappropriate for our purposes because it cannot directly be used to parametrize a surface in Boozer coordinates. This is because the toroidal angle used is the cylindrical angle ϕ, which does not correspond to the Boozer toroidal angle φ. It is not the choice of coordinate

system that is the issue as either cylindrical or Cartesian coordinates could describe a surface; we cannot directly use the VMEC representation due to the incompatibility of Boozer φ and cylindrical ϕ as described above.

3. Computing surfaces

Given a magnetic field with flux surfaces, Boozer coordinates (φ, θ) satisfy

$$\mathbf{B} = \nabla\Psi \times \nabla\theta + \iota\nabla\varphi \times \nabla\Psi, \tag{3.1}$$

where Ψ is the toroidal flux, derivatives are taken with respect to Cartesian coordinates (D'haeseleer *et al.* 2012), and we assume that $\iota \neq 0$. Note that (3.1) does not apply in regions in which the magnetic field lines are stochastic and fill a volume. In a vacuum, the magnetic field can be written as

$$\mathbf{B} = G\nabla\varphi, \tag{3.2}$$

where G is a constant. This is because $\nabla \times \mathbf{B} = 0$ and $\nabla \cdot \mathbf{B} = 0$ implies that $\mathbf{B} = \nabla V$ for some potential $V = G\varphi$. Taking the dot product of both sides of (3.1) and (3.2) with each other and dividing by G, we obtain

$$\nabla\Psi \cdot \nabla\theta \times \nabla\varphi = \frac{B^2}{G}, \tag{3.3}$$

where the field strength is given by $B := \|\mathbf{B}\|$. Using (3.3) with the following dual relations (Boozer 2005),

$$\nabla\varphi \times \nabla\Psi \frac{G}{B^2} = \frac{\partial \boldsymbol{\Sigma}}{\partial\theta}, \quad \nabla\Psi \times \nabla\theta \frac{G}{B^2} = \frac{\partial \boldsymbol{\Sigma}}{\partial\varphi}, \quad \text{and } \nabla\theta \times \nabla\varphi \frac{G}{B^2} = \frac{\partial \boldsymbol{\Sigma}}{\partial\Psi},$$

we conclude that surfaces represented in Boozer angles must satisfy

$$G\mathbf{B} - \|\mathbf{B}\|^2 \left(\frac{\partial \boldsymbol{\Sigma}}{\partial\varphi} + \iota \frac{\partial \boldsymbol{\Sigma}}{\partial\theta} \right) = 0, \tag{3.4a}$$

$$V(\boldsymbol{\Sigma}) - V_{\text{target}} = 0. \tag{3.4b}$$

Here, $\boldsymbol{B}$ is the magnetic field on $\boldsymbol{\Sigma}$. In (3.4b), $V(\boldsymbol{\Sigma})$ is the volume enclosed by $\boldsymbol{\Sigma}$ and $V_{\text{target}} \in \mathbb{R}\backslash\{0\}$ is a given target volume. The label constraint (3.4b) is necessary to make (3.4) a closed system. For the volume, we use the formula

$$V(\boldsymbol{\Sigma}) := \frac{1}{3} \int_0^1 \int_0^1 \boldsymbol{\Sigma} \cdot \left(\frac{\partial \boldsymbol{\Sigma}}{\partial\varphi} \times \frac{\partial \boldsymbol{\Sigma}}{\partial\theta} \right) \, d\varphi \, d\theta. \tag{3.5}$$

The unknowns in (3.4) are $G \in \mathbb{R}$, the surface $\boldsymbol{\Sigma}$, and the rotational transform ι. The solution to (3.4) is a magnetic surface for the field $\mathbf{B}$ that is parametrized in Boozer angles. In principle, we know that $G = n_{\text{fp}}(\mu_0/\pi) \sum_{k=1}^{N_c} I_k$ for an exact stellarator symmetric magnetic surface, where I_k is the current in the kth coil. However, we include G as an unknown since this may not be true for finite dimensional approximations of (3.4). As illustration of our parameterization, we show a surface that is everywhere tangent to a magnetic field generated by a subset of the NCSX coils in Figure 1, as well as curves of constant Boozer angle φ, θ on that surface. The Boozer angles φ, θ increase, respectively, the long and short way around the toroidal surface.

In the next sections, we discuss the nonlinear system resulting from replacing $\boldsymbol{\Sigma}$ with the finite-dimensional surface $\boldsymbol{\Sigma}_s(\varphi, \theta)$ from section 2 and enforcing (3.4) on a grid of collocation points. Depending on the number of degrees of freedom in the parametrization

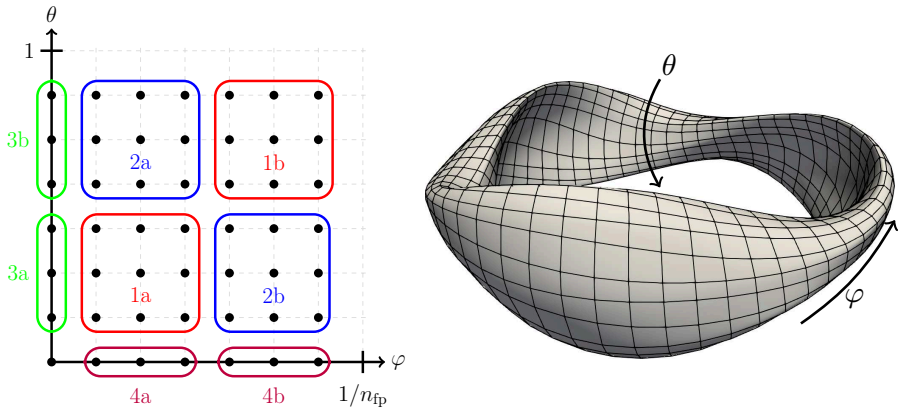

FIGURE 1. *Left:* Collocation points on a surface with $n_{\text{tor}} = m_{\text{pol}} = 3$, grouped by color. When stellarator symmetry is assumed, we choose one green, blue, purple, and red grouping of collocation points. Choosing points from two groups of the same color results in redundant equations. The y and z components of the residual at the origin are always included as the x component is zero due to stellarator symmetry. *Right:* Lines of constant Boozer angles φ, θ on a magnetic surface of the field generated by a subset of the NCSX coils (section 6 of Giuliani *et al.* (2020)). Arrows indicate increasing φ and θ, which respectively follow the long and short way around the torus. Note that the toroidal Boozer angle is not equal to the cylindrical toroidal angle.

and the number of collocation points, we either obtain a balanced system where the number of degrees of freedom is the same as the number of nonlinear equations, or an overdetermined system where the number of equations is larger than the number of unknowns. The derived PDE holds only in regions of the magnetic field where Boozer coordinates can be constructed. Despite this, we show an approach to nevertheless compute surfaces in regions of the magnetic field where nested flux surfaces do not exist.

3.1. *Balanced system*

We obtain a balanced system by using a parametrization $\boldsymbol{\Sigma}_s$ of $\boldsymbol{\Sigma}$ with unknowns $\mathbf{s} \in \mathbb{R}^{n_s}$, requiring (3.4) to be satisfied at the same number of collocation points, i.e.,

$$\begin{aligned} \mathbf{r}_{i,j}(\mathbf{s}) := G\mathbf{B}_{i,j} - \|\mathbf{B}_{i,j}\|^2 \left(\frac{\partial \boldsymbol{\Sigma}_{s,i,j}}{\partial \varphi} + \iota \frac{\partial \boldsymbol{\Sigma}_{s,i,j}}{\partial \theta} \right) &= 0, \ \forall (i,j) \in \mathcal{C} \\ V(\boldsymbol{\Sigma}_s) - V_{\text{target}} &= 0, \end{aligned} \tag{3.6}$$

where $\mathcal{C}$ is the set of indices of collocation points on a regular grid, $\boldsymbol{\Sigma}_{s,i,j} = \boldsymbol{\Sigma}_s(\varphi_{i,j}, \theta_{i,j})$, and $V(\boldsymbol{\Sigma}_s)$ is the volume enclosed by the surface, computed as

$$V(\boldsymbol{\Sigma}_s) = \frac{1}{3n'_\varphi n'_\theta} \sum_i \sum_j \boldsymbol{\Sigma}_{s,i,j} \cdot \left(\frac{\partial \boldsymbol{\Sigma}_{s,i,j}}{\partial \varphi} \times \frac{\partial \boldsymbol{\Sigma}_{s,i,j}}{\partial \theta} \right) \text{ for } (i,j) \in \mathcal{C}', \tag{3.7}$$

where $\mathcal{C}'$ is a set of quadrature points on a regular $n'_\varphi \times n'_\theta$ grid that is possibly different from the ones used by the collocation method $\mathcal{C}$ to ensure an accurate computation of the volume enclosed by the surface. The equations (3.6) require that the residual at a set of collocation points $(\varphi_{i,j}, \theta_{i,j})$ be zero. We place a grid of $(2n_{\text{tor}} + 1) \times (2m_{\text{pol}} + 1)$ equispaced collocation points on $[0, 1/n_{\text{nfp}}) \times [0, 1)$ and refer to surfaces that satisfy (3.6) at each point on that grid as BoozerExact surfaces.

Assuming stellarator symmetry, a number of these points are redundant and we

choose a subset of them. First, the x-component of the residual at the quadrature point associated to the origin, $i = j = 0$, or $\varphi = \theta = 0$, is always zero for all surfaces, so this equation can be ignored. To see why this is the case, note that the x-component of stellarator symmetric magnetic fields is zero when $\varphi = \theta = 0$. Similarly, differentiating the first equation in (2.1) with respect to φ or θ shows that the x-component of the surface tangents are zero when $\varphi = \theta = 0$. Second, if we include $(\varphi_{i,j}, \theta_{i,j})$ then the collocation point associated to $(-\varphi_{i,j}, -\theta_{i,j})$ is unnecessary. This is because the residual evaluated at $(\varphi_{i,j}, \theta_{i,j})$ is the residual at $(-\varphi_{i,j}, -\theta_{i,j})$ subject to a reflection. This is illustrated on the left in Figure 1, where the collocation points are grouped by color. We are free to choose one green, blue, purple, and red grouping. Choosing points from two groups of the same color would result in redundant equations. In all examples, we place collocation points on $[0, 1/2n_{\mathrm{fp}}) \times [0, 1)$ from zones 1a, 2a, 3a, and 4a, resulting in, respectively, $n_\varphi = n_{\mathrm{tor}} + 1$ and $m_\theta = 2m_{\mathrm{pol}} + 1$ quadrature points in the toroidal and poloidal directions (before removing redundant quadrature points in group 3b and at the origin).

After removing the redundant collocation points, the number of equations that must be satisfied is

$$3[2n_{\mathrm{tor}}m_{\mathrm{pol}} + m_{\mathrm{pol}} + n_{\mathrm{tor}} + 1] - 1 + 1 = 3[2n_{\mathrm{tor}}m_{\mathrm{pol}} + m_{\mathrm{pol}} + n_{\mathrm{tor}} + 1]. \tag{3.8}$$

This is the same as n_s+2, the number of geometric degrees of freedom in the parametrization n_s with ι and G. Since this is a balanced system of nonlinear equations, we can use Newton's method to solve it.

3.2. *Overdetermined system*

As we will show in numerical examples, the above balanced approach can fail to converge when the magnetic field lines have chaotic regions and islands. For such problems, we propose an alternative approach in which we relax the requirement that the residual be zero everywhere. Instead, we compute a surface that solves (3.4) in a least-squares sense. That is, it minimizes

$$f(\mathbf{s}) := \frac{1}{2}\sum_{i,j} \|\boldsymbol{r}_{i,j}(\boldsymbol{s})\|^2 + \frac{1}{2}w\left(V(\boldsymbol{\Sigma}_s) - V_{\mathrm{target}}\right)^2. \tag{3.9}$$

Here, the nonlinear least squares objective contains the sum of squares of (3.4) at a number of collocation points with a penalty term with weight $w > 0$ to approximately enforce the surface volume condition (3.4b). The first-order optimality condition for a minimizer of f is

$$\nabla_{\mathbf{s}} f = 0 \tag{3.10}$$

and the optimization problem can be solved using a Levenberg-Marquardt or Newton's method. In numerical experiments, we typically first compute a solution with the Levenberg-Marquardt method, which is known to work robustly for nonlinear least squares problems (Nocedal & Wright 2006). If the stationarity condition (3.10) is not satisfied to sufficient precision, then we further improve the solution using a few steps of Newton's method. We refer to surfaces that satisfy (3.10) as BoozerLS ("Boozer Least Squares") surfaces. Note that the BoozerLS surface has similarities with the quadratic flux minimizing (QFM) surfaces of Dewar *et al.* (2010). QFM surfaces are obtained by finding surfaces so that $(\boldsymbol{B}\cdot\boldsymbol{n})^2$ is minimized on them. This results in a surfaces that are approximately tangential to the magnetic field. Surfaces that satisfy (3.4a) have the stricter requirement that the magnetic field must be aligned in the direction $\frac{\partial \boldsymbol{\Sigma}}{\partial \varphi} + \iota \frac{\partial \boldsymbol{\Sigma}}{\partial \theta}$ and not just any tangential direction. In all examples, we place equispaced collocation

points on $[0, 1/2n_{\rm fp}) \times [0, 1)$, resulting in respectively $n_\varphi = n_{\rm tor} + 6$ and $m_\theta = 2m_{\rm pol} + 6$ quadrature points in the toroidal and poloidal directions.

3.3. *Example: Surfaces of an NCSX-like magnetic field*

In this section, we use a challenging magnetic field to study the behaviour of the BoozerExact and BoozerLS approaches for computing surfaces. Specifically, we use the 18 nonplanar coils and currents, corresponding to the C09R00 boundary configuration, of the national compact stellarator experiment (NCSX) to generate a vacuum magnetic field. The original NCSX design included planar toroidal and poloidal field coils that we do not use here. We have also previously considered this set up in section 6 of Giuliani *et al.* (2020). The field is challenging as it does not have nested flux surfaces everywhere, and presents an island chain in the neighborhood of the low-order rational $\iota = 3/7$. The magnetic field $\mathbf{B}$ on the surface $\boldsymbol{\Sigma}$ in (3.4a) is evaluated using the Biot-Savart law since we are working with vacuum fields. Newton's method requires the Jacobian of (3.6) or (3.10), which are functions of the magnetic field and its first spatial derivatives for (3.6), and of its first and second spatial derivatives for (3.10). Since the field is computed with the Biot-Savart law, we have explicit formulas to compute these quantities.

Due to the local convergence properties of the Newton and Levenberg-Marquardt algorithms to compute surfaces, they require initializations that should not be too far from the target surface. We use a continuation approach, i.e., we first compute a surface close to the axis, i.e., with a small value of $V_{\rm target}$. For that purpose, we generate a toroidal surface centered on the magnetic axis with a fixed minor radius. This does not correspond to a magnetic flux surface. However, in the neighborhood of the magnetic axis it is typically close enough for Newton's method to converge. We then use previously computed surfaces as initialization for finding the surface for the next larger $V_{\rm target}$.

In Figure 2a, we show cross sections of surfaces obtained with the BoozerExact approach, i.e., using the same number of unknowns as equations. In particular, we solve (3.6) with $m_{\rm pol} = n_{\rm tor} = 11$. As can be seen, when the Poincaré plot suggests the existence of magnetic surfaces, the BoozerExact approach converges to surfaces that align well with the Poincaré plot. In chaotic regions or regions containing magnetic islands, the approach may not converge, or it may find self-intersecting surfaces (as shown in red in the figure).

We next show that the BoozerLS approach is more robust in the presence of islands and chaotic regions. We again use $m_{\rm pol} = n_{\rm tor} = 11$ and the volume constraint penalty $w = 1000$, which results in a small relative error in the volume contained by the surface, ranging from 10^{-9} to 10^{-5}. The resulting cross sections are shown in Figure 2b. It can be seen that when the Poincaré plot indicates the existence of good surfaces, the BoozerLS surface cross sections visually coincide with the Poincaré plot and with the cross sections obtained with the balanced BoozerExact approach. In other regions, the BoozerLS approach results in layered smooth surfaces that do not have self-intersections (green curves). BoozerLS surfaces can either be computed using a standard continuation procedure, or alternatively, using possibly self-intersecting BoozerExact surfaces as initial guesses.

Finally, we numerically show that BoozerExact surfaces, i.e., solutions to (3.6), are not necessarily unique in the presence of chaos. This is more of theoretical interest than a fundamental problem as this is still a practical approach for computing magnetic surfaces when they exist. In addition, we will show in section 4 that it also can be used to optimize coils for quasi-symmetry. In Figure 2c, we show two BoozerExact surface cross sections with the same target volume enclosed, $V_{\rm target} = 1$. The red surfaces in Figure 2a and c are the same. The cyan cross section in Figure 2c is a BoozerExact surface initialized from the cyan BoozerLS surface in Figure 2b. It is still possible that the BoozerLS solver

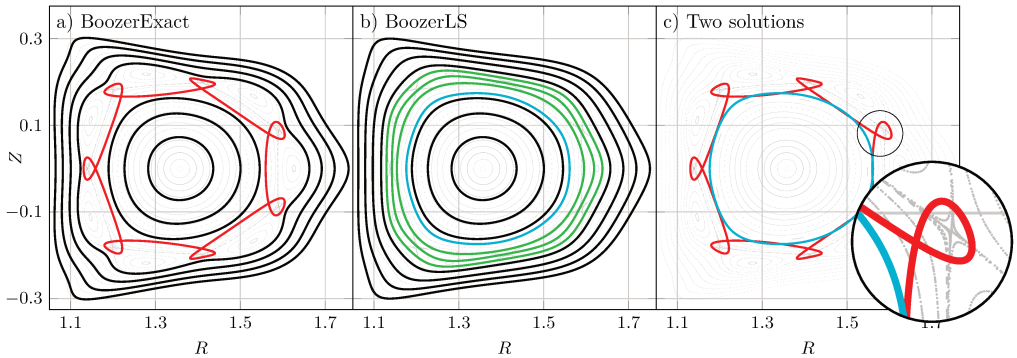

FIGURE 2. Poincaré plots (in grey) and surface cross sections for the field generated from a set of NCSX modular coils. at $\phi = \pi/3$. In a), we show cross sections of surfaces computed with the BoozerExact approach (in red and black). The red cross section is self-intersecting as it is close to an island chain. In b) we show cross section obtained with the BoozerLS algorithm (in black, green and cyan), which results in nested surfaces that do not self-intersect in the presence of islands. The red and cyan cross sections in a) and b) have the same target surface label, but are computed using BoozerExact and BoozerLS formulations, respectively. Green cross sections correspond to additional surfaces not shown in a). In c), we show two BoozerExact surfaces that have the same enclosed surface volume.

does not converge or for a self-intersecting BoozerLS surface to be found. However we find that this happens less frequently than it does for BoozerExact surfaces.

To further study the BoozerLS method in the presence of islands or chaotic fields, we examine the behavior of the residual as the number of Fourier modes (n_{tor}, m_{pol}) and collocation points used in the parametrization of the surface increases (Figure 3). As the number of degrees of freedom and collocation points in $\boldsymbol{\Sigma}_s$ increase, the residual decreases. However, this is only true for surfaces that do not pass through island chains. In the neighborhood of the island chain, we observe that the residual does not decrease with increasing surface complexity and appears to stagnate. This hints at an approach to optimize for magnetic fields without islands: adding the square of the residual at the collocation points of the BoozerLS surfaces to the objective function. This approach is not used in the coil optimization discussed next, which relies on BoozerExact surfaces. However, it will be the subject of future work.

4. Optimization for quasi-axisymmetry on surfaces

To target quasi-symmetry on magnetic surfaces within an optimization problem, we need a measure of quasi-symmetry on a given surface. For that purpose, we recall that the field strength $B := \|\mathbf{B}\|$ of a quasi-axisymmetric field on a magnetic surface in Boozer coordinates only depends on the angle θ. This motivates us to consider a decomposition into a quasi-axisymmetric and non-quasi-axisymmetric component

$$B(\varphi, \theta) = B_{\text{QS}}(\theta) + B_{\text{non-QS}}(\varphi, \theta).$$

This decomposition is not unique, and we choose $B_{\text{QS}}(\theta)$ such that it is closest to the field strength B when measured using the L_2-norm $(\int_{\boldsymbol{\Sigma}_s} f^2 \, dS)^{1/2}$, that is,

$$B_{\text{QS}}(\theta) = \frac{\int_0^{1/n_{\text{fp}}} B(\boldsymbol{\Sigma}_s(\varphi, \theta)) \, \|\frac{\partial \boldsymbol{\Sigma}_s}{\partial \varphi} \times \frac{\partial \boldsymbol{\Sigma}_s}{\partial \theta}\| \, d\varphi}{\int_0^{1/n_{\text{fp}}} \|\frac{\partial \boldsymbol{\Sigma}_s}{\partial \varphi} \times \frac{\partial \boldsymbol{\Sigma}_s}{\partial \theta}\| \, d\varphi}.$$

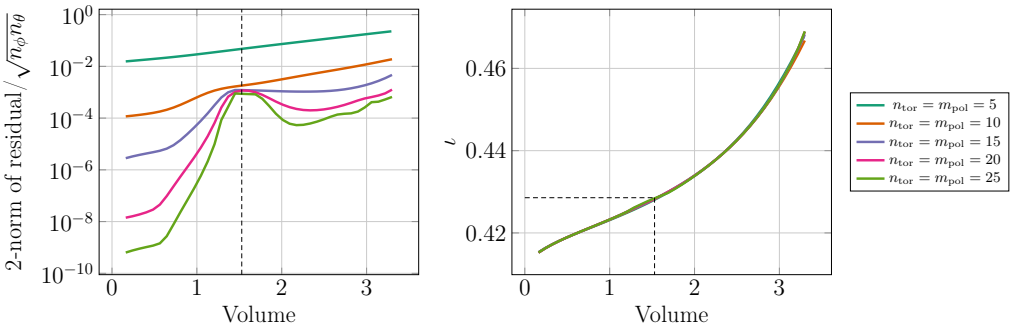

FIGURE 3. BoozerLS residual in the NCSX-lite configuration with increasing number of Fourier modes used to represent the surface. First figure shows the BoozerLS residuals as a function of the surface label. We observe that it stagnates in the neighborhood of $\iota = 3/7$, which corresponds to an island chain. Second figure shows the rotational transform with respect to surface label for different resolutions. The vertical dashed line approximately indicates the volume for which $\iota = 3/7$ is achieved.

Evaluating these quantities is straightforward because our surfaces are already parametrized in Boozer coordinates.

4.1. *Optimization problem formulation*

Now that we are able to compute surfaces and a measure for quasi-symmetry, the goal is to formulate and solve optimization problems for coils generating magnetic fields with good quasi-symmetry properties on a large number of toroidal surfaces. We seek a design with N_c independent modular coils, to which stellarator and rotational symmetries are applied. After application of these symmetries, the stellarator is made up of $2n_{\text{fp}}N_c$ modular coils. Each coil is parametrized with a current and n_c geometric degrees of freedom using a Fourier basis as in Giuliani *et al.* (2020). All coil degrees of freedom and their currents are summarized into a vector $\mathbf{c} \in \mathbb{R}^{N_c n_c + N_c}$. Our guiding principle for designing the coils is to target quasi-symmetry of the magnetic field induced by these coils on N_s surfaces. Each surface uses the parametrization from Section 2 with n_s geometric degrees of freedom. The resulting surface Fourier coefficients, ι, and G are summarized into a vector $\mathbf{s} \in \mathbb{R}^{N_s n_s + 2N_s}$. In the optimization formulation below, we use BoozerExact surfaces (Section 3.1). Finally, the optimization also includes various (inequality) constraints for the coils that are motivated by engineering considerations, which are critical in terms of the feasibility and cost-efficiency of a design (Strykowsky *et al.* 2009; Neilson *et al.* 2010; Klinger *et al.* 2013). These coil constraints are typically not difficult to compute as they explicitly depend on the coil degrees of freedom. This is in contrast to the terms that depend on the surfaces, which have an implicit dependence

on the coil degrees of freedom. The optimization problem is

$$\min_{\mathbf{c}\in\mathbb{R}^{N_c n_c + N_c},\, \mathbf{s}\in\mathbb{R}^{N_s n_s + 2N_s}} \hat{J}(\mathbf{c},\mathbf{s}) \tag{4.1a}$$

$$\text{subject to (3.6) on each of the } N_s \text{ surfaces,} \tag{4.1b}$$

$$\frac{1}{N_s}\sum_{k=1}^{N_s} \iota_k = \bar{\iota}, \tag{4.1c}$$

$$\Psi = \Psi_0, \tag{4.1d}$$

$$R_{\text{major}} = R_0, \tag{4.1e}$$

$$\sum_{i=1}^{N_c} L_c^{(i)} \leqslant L_{\max}, \tag{4.1f}$$

$$\kappa_i \leqslant \kappa_{\max},\ i = 1,\dots,N_c, \tag{4.1g}$$

$$\frac{1}{L_c^{(i)}}\int_{\boldsymbol{\Gamma}^{(i)}} \kappa_i^2\ dl \leqslant \kappa_{\text{msc}},\ i = 1,\dots,N_c, \tag{4.1h}$$

$$\|\boldsymbol{\Gamma}^{(i)} - \boldsymbol{\Gamma}^{(j)}\| \geqslant d_{\min} \quad \text{for } i \neq j, \tag{4.1i}$$

where the objective function is the average (normalized) non quasi-axisymmetry on the surfaces $\boldsymbol{\Sigma}_{s,k}$ $k = 1,\dots,N_s$:

$$\hat{J}(\mathbf{c},\mathbf{s}) = \frac{1}{N_s}\sum_{k=1}^{N_s} \frac{\int_{\boldsymbol{\Sigma}_{s,k}} B_{\text{non-QS},k}(\mathbf{c},\mathbf{s})^2\ dS}{\int_{\boldsymbol{\Sigma}_{s,k}} B_{\text{QS},k}(\mathbf{c},\mathbf{s})^2\ dS}. \tag{4.2}$$

The above objective is the average ratio of the squared 2-norm of the field magnitude's non-quasisymmetric and quasisymmetric components. The objective is subject to the surface constraints (4.1b)–(4.1e) and the coil constraints (4.1f)–(4.1i). In particular, (4.1c) ensures that the average rotational transform across the surfaces is $\bar{\iota}$. (4.1d) fixes the toroidal flux on the innermost surface to a given value Ψ_0 and prevents the currents from going to zero. (4.1e) fixes the major radius on the innermost surface to a given R_0 and prevents the length scale of the stellarator from changing. (4.1f) prevents the sum of the independent modular coil lengths $\sum_{i=1}^{N_c} L_c^{(i)}(\mathbf{c})$ from exceeding a given value $L_{\max} > 0$. (4.1g) and (4.1h), respectively, prevent the curvature and mean squared curvature on each coil from exceeding values $\kappa_{\max}$ and κ_{msc}. Finally, (4.1i) ensures that the coils stay at least $d_{\min} > 0$ away from one another.

4.2. *Computational aspects*

Our computational approach to finding minimizers for (4.1) uses a gradient-based descent algorithm to minimize the objective. This requires appropriate gradients of the objective, which take into account the equality PDE-constraints (4.1b) describing the surfaces, as well as the remaining equality and inequality constraints (4.1c)–(4.1i). To enforce the PDE-constraints (4.1b), we use the method of Lagrange multipliers (Tröltzsch 2010; Giuliani *et al.* 2020). For that purpose, we consider the objective as a function just of the coil degrees of freedom, i.e., $J(\mathbf{c}) := \hat{J}(\mathbf{c},\mathbf{s}(\mathbf{c}))$ by considering the surfaces as functions of the coils, i.e., $\mathbf{s} = \mathbf{s}(\mathbf{c})$ implicitly defined through the solution of (3.6). Computing the gradient of J with respect to $\mathbf{c}$ can be done efficiently by introducing Lagrange multipliers and using the adjoint method. Computing gradients then requires solving two systems of equations per surface, both with $n_s + 2$ equations: the nonlinear system of equations (3.6) and a linear system for the adjoint variable. This method

compares favorably to finite difference approaches for computing gradients, where the number of times (3.6) must be solved scales with n_c, the number of degrees of freedom used to represent the coils.

The constraints (4.1c)–(4.1i) are enforced approximately using a penalty method, with penalties that only become active when the constraint is violated. We next detail the penalty terms used, and discuss a practical strategy to choose and adjust the weights for these penalties at the end of this section.

For the penalty on the sum of the independent modular coil lengths, we use the standard quadratic penalty

$$\frac{1}{2}\left[\max\left(0, \sum_{i=1}^{N_c} L_c^{(i)}(\mathbf{c}) - L_{\max}\right)\right]^2 .$$

The penalty on the maximum curvature for coil $\boldsymbol{\Gamma}^{(i)}$ is

$$\sum_{i=1}^{N_c} \frac{1}{2} \int_{\boldsymbol{\Gamma}^{(i)}} \left[\max\left(0, \kappa_i(\mathbf{c}) - \kappa_{\max}\right)\right]^2 \, dl_i,$$

and, as in Giuliani *et al.* (2020), the penalty on the minimum pairwise coil distance is

$$\frac{1}{2} \sum_{i \neq j} \int_{\boldsymbol{\Gamma}^{(i)}} \int_{\boldsymbol{\Gamma}^{(j)}} \left[\max(0, d_{\min} - \|\boldsymbol{\Gamma}^{(i)} - \boldsymbol{\Gamma}^{(j)}\|_2)\right]^2 \, dl_j \, dl_i. \tag{4.3}$$

For the equality constraints, we add quadratic penalties to the objective. For instance, to prevent the currents in the coils from becoming zero, for the toroidal flux Ψ of the innermost surface we use

$$\frac{1}{2}\left(\Psi(\mathbf{c}, \mathbf{s}) - \Psi_0\right)^2 .$$

Here, the toroidal flux through a toroidal cross section of constant Boozer angle $\varphi = 0$, $S_{\varphi=0}$, of a surface is given by

$$\Psi(\mathbf{c}, \mathbf{s}) = \int_{S_{\varphi=0}} \mathbf{B} \cdot \mathbf{n} \, dS = \int_0^1 \mathbf{A} \cdot \frac{\partial}{\partial \theta} \boldsymbol{\Sigma}_s(0, \theta) \, d\theta, \tag{4.4}$$

where $\mathbf{A}$ is the vector potential, such that $\mathbf{B} = \nabla \times \mathbf{A}$, and where we have used Stokes' theorem to turn the surface integral for the flux into a line integral. Additionally, we include a penalty term on the major radius R_{major} of the innermost surface:

$$\frac{1}{2}\left(R_{\text{major}}(\mathbf{s}) - R_0\right)^2 .$$

This provides a characteristic length scale to the stellarator configurations we optimize, and maintains the aspect ratio of the surfaces in the volume. Details of how we compute this quantity are given in Appendix B. To constrain the mean rotational transform of the surfaces to be close to a given value $\bar{\iota}$, we use the penalty term

$$\frac{1}{2}\left(\frac{1}{N_s} \sum_{k=1}^{N_s} \iota_k(\mathbf{s}) - \bar{\iota}\right)^2 ,$$

where ι_k is the rotational transform on the kth surface. While here we target the average rotational transform, other choices are possible in our framework, such as constraining any number of surface rotational transforms to a given value or interval. Finally, we

include a penalty that favors coils with an approximately uniform incremental arclength

$$\sum_{i=1}^{N_c} \int_0^1 (\|\boldsymbol{\Gamma}^{(i)\prime}\|_2 - t^{(i)})^2 \, d\xi,$$

with $t^{(i)} = \int_0^1 \|\boldsymbol{\Gamma}^{(i)\prime}\|_2 \, d\xi$, where $\xi \in [0,1)$ is the angle parameterization of the coil. This prevents the incremental arclengths from getting close to zero, where the length penalty term is nondifferentiable. Note that it might not be possible to attain a uniform incremental arclength at all points on the coil. However, this is not necessary as the purpose of this penalty is to avoid the above issue. Numerical experiments confirm that this term does not substantially affect the value of the objective over a wide range of penalty weights.

We compute analytical gradients of the (reduced) objective $J(\mathbf{c})$, and use the BFGS method (Nocedal & Wright 2006) with linesearch to find a minimizer. Compared to steepest descent, BFGS has the advantage that it approximates second-order information of the objective, which typically results in faster convergence or, in some cases, avoids that the iteration progress stalls.

To choose appropriate penalty weights, we follow a simple strategy that appropriately increases these weights throughout the iteration. Namely, if any of the constraints are not satisfied to within a 0.1% relative tolerance after 3,000 iterations of BFGS, we increase the weights on the associated penalty terms by a factor of 10. Additionally, we allow for a maximum number of 45,000 BFGS iterations. However, since we choose a judicious set of initial weights informed from previous runs, weights are increased rarely.

4.3. *Coil optimization for precise quasi-axisymmetry on surfaces*

In this example, we revisit a coil set obtained in (Wechsung *et al.* 2022) with the second stage of the standard two-stage optimization method, where the goal was to produce the precise QA magnetic field presented in Landreman & Paul (2022) using a realistic set of coils. In the first stage, that field was optimized for precise QA on a toroidal volume with aspect ratio 6, and the average rotational transform in the volume was 0.42. The stage II optimization found $N_c = 4$ coils with a minimum pairwise coil distance of $d_{\min} = 0.1\,\mathrm{m}$, and maximum curvature and mean squared curvature of $\kappa_{\max} = 5\,\mathrm{m}^{-1}$ and $\kappa_{\mathrm{msc}} = 5\,\mathrm{m}^{-2}$. Various coil designs are found where the the sum of the lengths of the 4 base coils is constrained to be below $L_{\max} = 18\,\mathrm{m}, 20\,\mathrm{m}, 22\,\mathrm{m}, 24\,\mathrm{m}$. Here, the major radius of the device is approximately 1 m. In the following optimization, we use the same values for $d_{\min}$, $\kappa_{\max}$, κ_{msc} as in the stage II optimization from Wechsung *et al.* (2022).

Building on the results from this two-stage coil design, we aim at improving the quasi-symmetry on multiple surfaces uniformly distributed throughout the volume using the optimization approach presented in the previous section. Since the fields resulting from the coils in Wechsung *et al.* (2022) appear to avoid islands, we use BoozerExact (rather than BoozerLS) surfaces in this optimization. Should islands appear in the vicinity of the surfaces over the course of the optimization, the BoozerExact surfaces might begin to self-intersect and no longer correspond to flux surfaces of the magnetic field that we are trying to optimize. However, as will be shown, this does not happen in the final coil designs. Generalizing the formulation (4.1) to BoozerLS surfaces, in order to handle magnetic fields with islands in a robust manner, and to optimize the coils to eliminate such islands, will be part of future work.

We first optimize coils using five BoozerExact surfaces uniformly spaced in terms of the surface label $\sqrt{V}$, which is proportional to the minor radius. We choose $m_{\mathrm{pol}} = n_{\mathrm{tor}} = 10$, and the outermost surface in the initial configuration has aspect ratio approximately

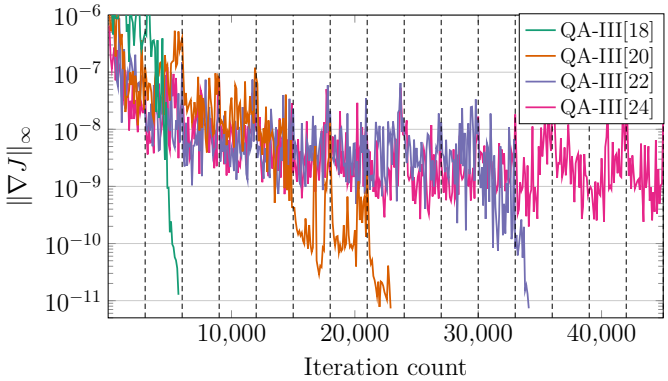

FIGURE 4. Norm of gradient in each iteration where five surfaces are used. Gradients at the initial configuration were $\mathcal{O}(10^{-4})$, but we truncate the y-axis in order to better observe convergence to optimality. Every 3,000 iterations (indicated by dotted vertical lines), the penalty weights are increased until the (in)equality constraints in (4.1) are satisfied to sufficient precision. Since a suitable set of initial weights were chosen, this does not occur frequently.

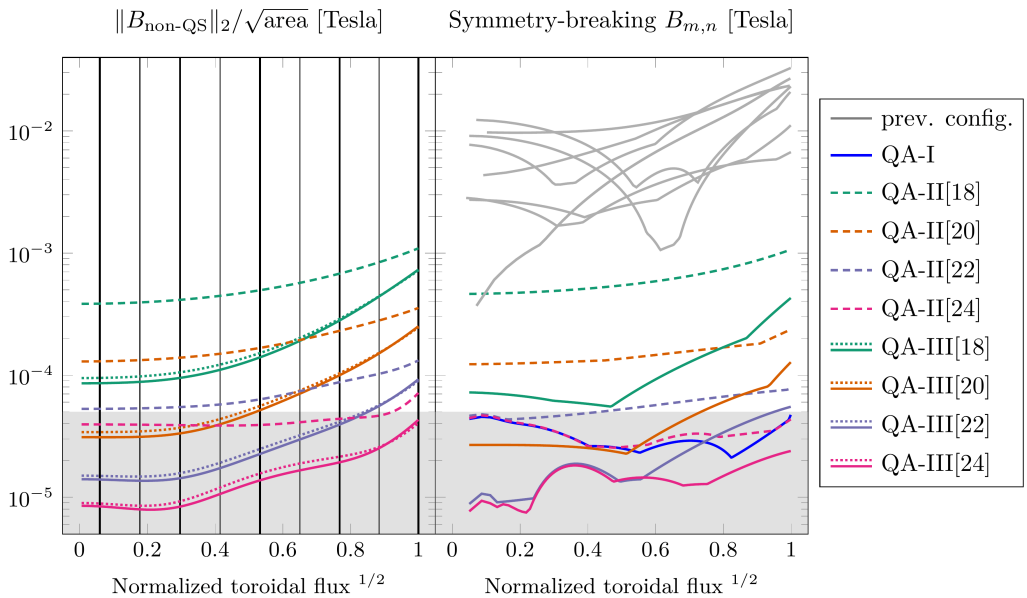

FIGURE 5. Two measures of the field's departure from quasi-symmetry: the ratio $\|B_{\text{non-QS}}\|/\sqrt{\text{area}}$ (left), and the maximum symmetry-breaking Fourier mode (right). In both plots, the magnetic field is normalized to be 1 T on the magnetic axis. The curves corresponding to the QA-III coils in the left figure have two line styles: dotted (·····), and solid (——), which correspond to configurations resulting from optimization on five and nine surfaces respectively. The maximum symmetry-breaking Fourier mode is calculated by computing a two-dimensional FFT of the field strength, and finding the maximum $|B_{m,n}|$ for $n \neq 0$, corresponding to the Fourier modes introducing variation in φ. The thick and thin vertical black lines correspond respectively to the five surfaces used in the initial optimization and the four surfaces added in the nine surface optimization run. The shaded gray region of the plot corresponds to field strengths below 50 µT, corresponding to the Earth's background magnetic field.

equal to 6, similar to the optimization in Landreman & Paul (2022). We set the target toroidal flux (equation (4.1*d*)) and major radius (equation (4.1*e*)) of the innermost surface to be the same as the one in the stage II coil set. In addition, we use the same coil engineering requirements (equations (4.1*f*)-(4.1*i*)) as the ones used in Wechsung *et al.* (2022). Each of the four base modular coils is described by a Fourier representation with 16 Fourier modes, which results in an optimization problem with 400 degrees of freedom corresponding to the geometry and coil currents.

Unlike in Landreman & Paul (2022), we do not directly fix the aspect ratio of the outermost surface. Instead, we fix the major radius of the innermost surface and the volume of the outermost surface. We empirically observe that this effectively limits the variation of the aspect ratio of the outermost surface over the course of the optimization. Once the coils are optimized on five surfaces, we increase the number of BoozerExact surfaces to nine by keeping the original five surfaces and introducing four new surfaces in between, and again run at most 45,000 BFGS iterations.

In the following description of the results, we will refer to the stage I precise QA magnetic field of Landreman & Paul (2022) as QA-I, the configurations obtained from the classical stage II coil optimization in Wechsung *et al.* (2022) for different maximum coil lengths as QA-II[$L_{\max}$], and the stage III coils obtained from solving the optimization problem described in this manuscript as QA-III[$L_{\max}$].

When optimizing on nine surfaces, all coil sets considered in this manuscript reached optimality, ($\|\nabla J\|_\infty < 10^{-10}$), within the maximum number of iterations and the gradient was reduced by at least six orders of magnitude. More iterations are required for longer coil lengths since the optimization problem becomes more ill-conditioned (see Figure 4). In Figure 5, we provide two similar measures of quasi-symmetry: the norm of the non-quasisymmetric component of the field magnitude and the maximum symmetry-breaking Fourier modes. Significant improvements in quasi-symmetry are observed after optimization on five surfaces. Additional minor improvement is observed when increasing from five to nine surfaces (Figure 5a). This suggests that for this particular setup, five surfaces were sufficient to optimize for QA throughout the volume and not just in the neighborhood of the surfaces.

In our implementation, we use MPI-based parallelism to launch N_s tasks that concurrently compute surfaces and adjoint variables, meaning that the time taken to evaluate the objective and its gradient does not substantially increase with the number of surfaces on which quasi-symmetry is optimized. Since five surfaces might not be sufficient in all configurations, the possibility of parallelizing the surface and adjoint solves, as available in our implementation, means that we do not have to compromise on the number of surfaces used. Our implementation relies heavily on the open-source stellarator optimization package SIMSOPT (Landreman *et al.* 2021) where many aspects of this work are contributed. We also use SIMSOPT's "optimizable" framework for efficient representation of our objective (appendix C). In the above example, a typical iteration of BFGS takes approximately two seconds on Intel Xeon Platinum 8268 processors when optimizing on five or nine surfaces.

In Figure 5b, we provide the magnitude of the maximum symmetry-breaking mode of the field magnitude and compare with eight previous configurations taken from figure 6a in (Landreman & Paul 2022). The shaded area of the figure indicates field magnitudes that are smaller than the Earth's background magnetic field (50 µT). We highlight that after optimization, the symmetry-breaking errors of the QA-III[24] coils improve not only on the QA-II[24] field, but even on the original QA-I field. This might have a number of explanations. For example, the dimension of the optimization problem here is 400 and thus just over three times larger than the dimension of the problem used to design QA-I,

$L_{\max}$	Maximum curvatures	Mean-squared curvatures	Coil-coil separation	Coil-surf. separation
QA-II[18]	$4.33, 4.42, 4.29, 4.47$	$4.56, 4.85, 4.90, 4.98$	0.132	0.271
QA-II[20]	$4.89, 4.84, 4.67, 5.00$	$4.60, 5.00, 5.00, 5.00$	0.113	0.287
QA-II[22]	$5.00, 5.00, 5.00, 5.00$	$4.17, 5.00, 5.00, 5.00$	0.099	0.303
QA-II[24]	$5.00, 5.00, 5.00, 5.00$	$3.99, 4.92, 5.00, 5.00$	0.099	0.307
QA-III[18]	$3.54, 4.07, 4.74, 4.36$	$5.00, 5.00, 5.00, 4.30$	0.108	0.239
QA-III[20]	$4.76, 5.00, 5.00, 4.97$	$5.00, 5.00, 5.00, 5.00$	0.099	0.272
QA-III[22]	$4.87, 5.00, 5.00, 5.00$	$5.00, 5.00, 5.00, 5.00$	0.099	0.289
QA-III[24]	$5.00, 5.00, 5.00, 5.00$	$5.00, 5.00, 5.00, 5.00$	0.099	0.297

TABLE 1. Comparison of geometric coil properties obtained from stage II and stage III optimization for different modular coil lengths. Stage III results are optimized using nine surfaces. Coil-to-surface distance is computed with respect to the surface used for coil design in the stage II rows, and the outermost surface on which quasi-symmetry is optimized for the stage III rows.

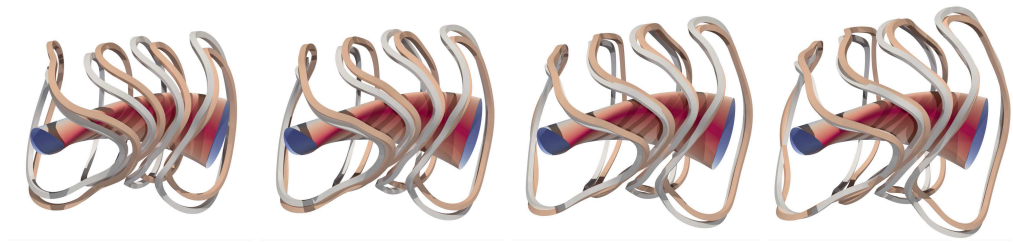

FIGURE 6. From left to right, inboard view of the $18\,\mathrm{m}, 20\,\mathrm{m}, 22\,\mathrm{m}, 24\,\mathrm{m}$ coils before and after optimization. Grey and gold correspond respectively to the stage II coils and stage III coils (nine surface optimization). The color on the surface corresponds to the magnetic field strength.

which is 120. This also might be because we use discretely exact gradients in this work, and all coils in the nine surface optimization reached optimality, to very close precision. This is in contrast to the QA-I field computations, where only finite difference gradients were used, thus potentially limiting the accuracy of gradients and as a result how much the objective could be decreased.

In the physics results discussed below, we now only analyze the coils optimized for QA using nine surfaces. Geometric properties of the final coils after optimization for QS on nine surfaces are provided in Table 1 and the coils are shown in Figure 6. The full coil set for the QA-III[24] configuration is shown in Figure 8 (left). As $L_{\max}$ increases, the coil-to-surface separation increases as coils can move further away from the outermost surface, thereby reducing the effects of the discrete nature of the coils and producing better quasi-symmetry in the volume. As the coil length increases, the maximum curvatures, mean-squared curvatures, and the coil-coil separation approach and attain $\kappa_{\max}$, κ_{msc}, $d_{\min}$, respectively. We observe that although the surfaces do not change substantially in the optimization (Figure 9), the coils do (Figure 6).

Cross sections of surfaces in the magnetic field after stage III optimization on nine surfaces and Poincaré plots are provided in Figure 9 (left), where there are no visible islands. The rotational transform is approximately constant and equal to $\bar{\iota}$ throughout the volume on which we are optimizing, varying between 0.415 and 0.428 for the different coil lengths. We also compare cross sections of surfaces in the stage II and stage III fields

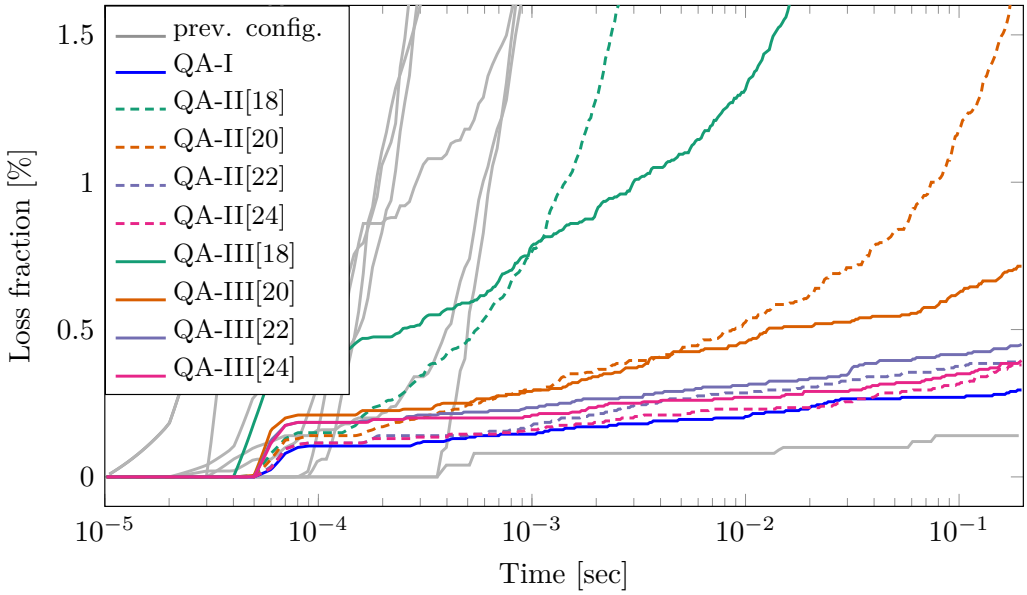

FIGURE 7. Losses of alpha particles spawned on the surface with normalized toroidal flux 0.25. We observe significant improvements in confinement for the shorter coils. For $L_{\max} = 18\,\mathrm{m}$, losses are reduced from 17.7% to 6.6%, for $L_{\max} = 20\,\mathrm{m}$, losses are reduced from 1.7% to 0.7%. The stage II and III coils with lengths $L_{\max} = 22\,\mathrm{m}, 24\,\mathrm{m}$ present comparable particle losses, despite significantly better quasi-symmetry for the stage III coil sets. The previous configurations, shown in grey, are detailed in Landreman & Paul (2022).

in Figure 9 (right). We observe that although the coils change substantially (Figure 6), the surfaces do not.

We compute two measures to quantify the confinement properties of our optimized magnetic configurations: alpha particle losses (Figure 7) and the thermal collisional transport magnitude $\epsilon_{\mathrm{eff}}^{3/2}$ (Figure 8 right). The alpha particle losses are calculated as follows. Each configuration is first scaled to match the minor radius and average field strength of the ARIES-CS reactor, 1.70 m and 5.9 T. Alpha particles are initialized isotropically in velocity space with 3.5 MeV energy, and uniformly on the flux surface with toroidal flux 1/4 that of the outermost computational surface, corresponding to half of the effective minor radius. The collisionless guiding-center trajectories are followed using the SIMPLE code (Albert *et al.* 2020*a*,*b*), and particles are considered lost if they exit the outermost surface. When considering the confinement of alpha particles, we see the most significant improvements for the two shorter coil sets. For QA-III[18], losses are reduced from 17.7% to 6.6%, for QA-III[20], losses are reduced from 1.7% to 0.7%. For the longer QA-III[22] and QA-III[24] coils, the (small) particle losses are comparable to those for QA-II[22] and QA-II[24] respectively, despite significantly better quasi-symmetry throughout the volume. We note that this imperfect correlation between quasi-symmetry and energetic particle confinement is consistent with the findings in Bader *et al.* (2021), and in Landreman & Paul (2022) where the QA+Well configuration had lower particle losses than the QA configuration despite worse quasi-symmetry. For $\epsilon_{\mathrm{eff}}^{3/2}$ (Nemov *et al.* 1999) we observe significant improvement compared to the coils of Wechsung *et al.* (2022), and our procedure is even able to improve the results from the target field of Landreman & Paul (2022).

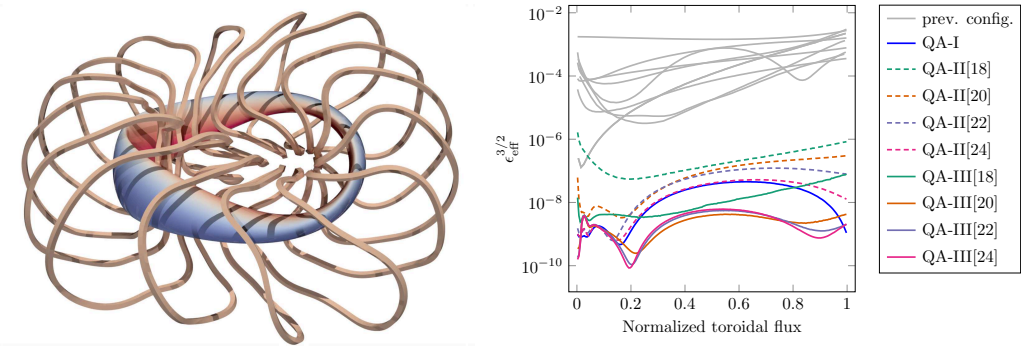

FIGURE 8. *Left*: The QA-III[24] stellarator and the outermost surface on which quasi-symmetry is optimized. *Right*: The thermal collisional transport magnitude $\epsilon_{\text{eff}}^{3/2}$ is significantly reduced; the QA-III[20], QA-III[22], and QA-III[24] configurations even improve on the QA-I configuration. The previous configurations, shown in grey, are detailed in Landreman & Paul (2022).

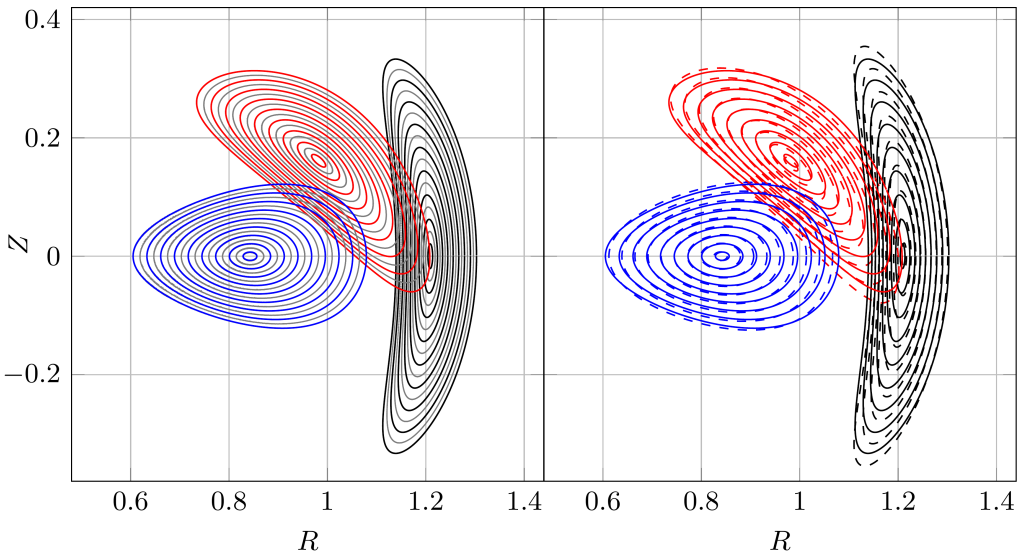

FIGURE 9. Shown on the left and right with solid lines (——) are cross sections of the nine surfaces used in the QA-III[18] coil set optimization at $\phi = 0, \pi/4, \pi/2$ in black, red and blue. Grey lines on the left correspond to Poincare plots. Shown on the right are cross sections of surfaces in the QA-II[18] field using dashed lines (- - -). The longer coil sets result in comparable cross sections, but the difference between stage II and III cross sections become smaller.

5. Conclusions

The main contributions of this work are twofold: the introduction of a novel approach for computing surfaces directly parametrized in Boozer coordinates (section 3), and the optimization of a vacuum magnetic field generated by coils for quasi-symmetry on a number of magnetic surfaces (section 4).

The method we call "BoozerExact surfaces" can be used to accurately compute magnetic surfaces when islands or chaos are not present. In contrast, we have shown that the method we call "BoozerLS surfaces" is a more robust approach to compute surfaces, even in regions where nested flux surfaces do not exist. In addition, the BoozerLS residual can effectively be used to identify regions where the assumption of nested flux surfaces is no longer valid.

In the second part of the article, we have shown that our new framework to compute magnetic surfaces can be used to optimize the geometry and currents of filamentary coils, with the goal of improving the quasi-symmetry of the magnetic field generated by these coils. At the moment, this approach requires an initial coil set that already produces a magnetic field with surfaces that can be parametrized in Boozer coordinates. In other words, it requires a "warm-start". In addition, the method only works in a robust manner if islands or chaotic regions do not appear over the course of the coil optimization. In section 4.3, we demonstrate that this technique is highly effective when these requirements are satisfied. In fact, for the QA-III[24] configuration, we found a coil design with quasi-symmetry that improved on the quasi-symmetry of the original QA-I configuration. This could be attributed to the higher dimensionality of our optimization problem, or our use of analytical gradients as opposed to those obtained using finite differences. For the shorter coil configuration QA-III[18], we are able to significantly improve performance compared to coils obtained from the standard two-stage approach: the thermal collisional transport magnitude $\epsilon_{\text{eff}}^{3/2}$ is reduced by more than an order of magnitude, and alpha particle losses are reduced from 17.7% to 6.6%. Future work will include using BoozerLS surfaces and their residual to optimize coils with less suitable initial coil configurations. We also will consider a stochastic version of this objective to find coils that are robust to manufacturing errors.

Availability of code and optimized configurations

The coil and surface parametrizations, Biot-Savart kernel, as well as the solver used to compute BoozerExact and BoozerLS surfaces were implemented in the SIMSOPT package (Landreman *et al.* 2021), available at:

`https://github.com/hiddenSymmetries/simsopt`.

Driver scripts for work are available at:

`https://github.com/andrewgiuliani/PySurfaceOpt`.

The stage III coils discussed in this work are available at the above repository as well.

Acknowledgements

The authors would like to thank the SIMSOPT development team. This work was supported by a grant from the Simons Foundation (560651). AG is partially supported by an NSERC (Natural Sciences and Engineering Research Council of Canada) postdoctoral fellowship. In addition, AC and FW are supported by the United States National Science Foundation under grant No. PHY-1820852, and AC is supported by the United States Department of Energy, Office of Fusion Energy Sciences, under grant No. DE-FG02-86ER53223.

REFERENCES

Albert, C G, Kasilov, S V & Kernbichler, W 2020*a* Accelerated methods for direct computation of fusion alpha particle losses within stellarator optimization. *Journal of Plasma Physics* **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. *Journal of Computational Physics* **403**, 109065.

Bader, A, Anderson, D T, Drevlak, M, Faber, B J, Hegna, C C, Henneberg, S, Landreman, M, Schmitt, J C, Suzuki, Y & Ware, A 2021 Modeling of energetic particle transport in optimized stellarators. *Nuclear Fusion* **61**, 116060.

Boozer, Allen H. 1981 Plasma equilibrium with rational magnetic surfaces. *The Physics of Fluids* **24** (11), 1999–2003, arXiv: https://aip.scitation.org/doi/pdf/10.1063/1.863297.

Boozer, Allen H. 2005 Physics of magnetically confined plasmas. *Rev. Mod. Phys.* **76**, 1071–1141.

Brown, T., Breslau, J., Gates, D., Pomphrey, N. & Zolfaghari, A. 2015 Engineering optimization of stellarator coils lead to improvements in device maintenance. In *2015 IEEE 26th Symposium on Fusion Engineering (SOFE)*, pp. 1–6.

Dewar, RL, Gibson, AM & Hudson, SR 2010 Unified theory of ghost and quadratic-flux-minimizing surfaces. *Journal of Plasma and Fusion Research SERIES* **9**.

D'haeseleer, William D, Hitchon, William NG, Callen, James D & Shohet, J Leon 2012 *Flux coordinates and magnetic field structure: a guide to a fundamental tool of plasma theory*. Springer Science & Business Media.

Drevlak, Michael 1998 Automated optimization of stellarator coils. *Fusion Technology* **33** (2), 106–117.

Garren, D. A. & Boozer, A. H. 1991 Existence of quasihelically symmetric stellarators. *Physics of Fluids B: Plasma Physics* **3** (10), 2822–2834.

Giuliani, Andrew, Wechsung, Florian, Cerfon, Antoine, Stadler, Georg & Landreman, Matt 2020 Single-stage gradient-based stellarator coil design: Optimization for near-axis quasi-symmetry, arXiv: 2010.02033.

Helander, Per 2014 Theory of plasma confinement in non-axisymmetric magnetic fields. *Reports on Progress in Physics* **77** (8), 087001.

Henneberg, Sophia A, Hudson, Stuart R, Pfefferlé, David & Helander, Per 2021 Combined plasma–coil optimization algorithms. *Journal of Plasma Physics* **87** (2).

Hirshman, S. P. & Whitson, J. C. 1983 Steepest-descent moment method for three-dimensional magnetohydrodynamic equilibria. *The Physics of Fluids* **26** (12), 3553–3568, arXiv: https://aip.scitation.org/doi/pdf/10.1063/1.864116.

Klinger, T., Baylard, C., Beidler, C.D., Boscary, J., Bosch, H.S., Dinklage, A., Hartmann, D., Helander, P., Massberg, H., Peacock, A., Pedersen, T.S., Rummel, T., Schauer, F., Wegener, L. & Wolf, R. 2013 Towards assembly completion and preparation of experimental campaigns of Wendelstein 7-X in the perspective of a path to a stellarator fusion power plant. *Fusion Engineering and Design* **88** (6), 461 – 465, proceedings of the 27th Symposium On Fusion Technology (SOFT-27); Liège, Belgium, September 24-28, 2012.

Kruger, Thomas G., Zhu, C., Bader, A., Anderson, D. T. & Singh, L. 2021 Constrained stellarator coil curvature optimization with FOCUS. *Journal of Plasma Physics* **87** (2), 175870201.

Landreman, Matt 2017 An improved current potential method for fast computation of stellarator coil shapes. *Nuclear Fusion* **57** (4), 046003.

Landreman, M, Medasani, B, Wechsung, F, Giuliani, A, Jorge, R & Zhu, C 2021 SIMSOPT: a flexible framework for stellarator optimization. *Journal of Open Source Software* **6**, 3525.

Landreman, Matt & Paul, Elizabeth 2022 Magnetic fields with precise quasisymmetry for plasma confinement. *Phys. Rev. Lett.* **128**, 035001.

Landreman, Matt & Sengupta, Wrick 2018 Direct construction of optimized stellarator shapes. Part 1. Theory in cylindrical coordinates. *Journal of Plasma Physics* **84** (6), 905840616.

Landreman, Matt & Sengupta, Wrick 2019 Constructing stellarators with quasisymmetry to high order. *Journal of Plasma Physics* **85** (6).

Landreman, Matt, Sengupta, Wrick & Plunk, Gabriel G. 2019 Direct construction of optimized stellarator shapes. Part 2. Numerical quasisymmetric solutions. *Journal of Plasma Physics* **85** (1), 905850103.

Merkel, P. 1987 Solution of stellarator boundary value problems with external currents. *Nuclear Fusion* **27** (5), 867–871.

Neilson, G. H., Gruber, C. O., Harris, J. H., Rej, D. J., Simmons, R. T. & Strykowsky, R. L. 2010 Lessons learned in risk management on NCSX. *IEEE Transactions on Plasma Science* **38** (3), 320–327.

Nemov, V V, Kasilov, S V, Kernbichler, W & Heyn, M F 1999 Evaluation of $1/\nu$ neoclassical transport in stellarators. *PoP* **6**, 4622.

Nocedal, Jorge & Wright, Stephen 2006 *Numerical optimization*. Springer Science & Business Media.

Paul, E.J., Landreman, M., Bader, A. & Dorland, W. 2018 An adjoint method for gradient-based optimization of stellarator coil shapes. *Nuclear Fusion* **58** (7), 076015.

Singh, Luquant, Kruger, T. G., Bader, A., Zhu, C., Hudson, S. R. & Anderson, D. T. 2020 Optimization of finite-build stellarator coils. *Journal of Plasma Physics* **86** (4), 905860404.

Spong, D A 2015 3d toroidal physics: Testing the boundaries of symmetry breaking. *Physics of Plasmas* **22**, 055062.

Strickler, Dennis J., Berry, Lee A. & Hirshman, Steven P. 2002 Designing coils for compact stellarators. *Fusion Science and Technology* **41** (2), 107–115.

Strickler, Dennis J., Hirshman, Steven P., Spong, Donald A., Cole, Michael J., Lyon, James F., Nelson, Bradley E., Williamson, David E. & Ware, Andrew S. 2004 Development of a robust quasi-poloidal compact stellarator. *Fusion Science and Technology* **45** (1), 15–26.

Strykowsky, R. L., Brown, T., Chrzanowski, J., Cole, M., Heitzenroeder, P., Neilson, G. H., Rej, D. & Viol, M. 2009 Engineering cost schedule lessons learned on NCSX. In *2009 23rd IEEE/NPSS Symposium on Fusion Engineering*, pp. 1–4.

Tröltzsch, Fredi 2010 *Optimal Control of Partial Differential Equations: Theory, Methods and Applications*, *Graduate Studies in Mathematics*, vol. 112. American Mathematical Society.

Wechsung, Florian, Giuliani, Andrew, Landreman, Matt, Cerfon, Antoine & Stadler, Georg 2021 Single-stage gradient-based stellarator coil design: stochastic optimization, arXiv: 2106.12137.

Wechsung, Florian, Landreman, Matt, Giuliani, Andrew, Cerfon, Antoine & Stadler, Georg 2022 Precise stellarator quasi-symmetry can be achieved with electromagnetic coils. *Proceedings of the National Academy of Sciences (to appear)* .

Yu, Guodong, Feng, Zhichen, Fu, Guoyong & others 2022 Existence of an optimized stellarator with simple coils. *arXiv preprint arXiv:2201.12547* .

Zhu, Caoxiang, Hudson, Stuart R, Song, Yuntao & Wan, Yuanxi 2017 New method to design stellarator coils without the winding surface. *Nuclear Fusion* **58** (1), 016008.

A. Properties of the surface discretization

We give a brief derivation of the surface discretization described in Section 2. First, note that any periodic smooth function f in one dimension can be approximated arbitrarily well by functions of the form

$$f_{n_{\text{tor}}}(\varphi) = \sum_{i=1}^{2n_{\text{tor}}+1} a_i v_i(\varphi), v_i \in V^{n_{\text{tor}}}$$

where $V^{n_{\text{tor}}} = \{1, \cos(2\pi\varphi), \ldots, \cos(n_{\text{tor}}2\pi\varphi), \sin(2\pi\varphi), \ldots, \sin(n_{\text{tor}}2\pi\varphi)\}$. Hence, by considering the tensor product of this Fourier basis, we obtain that any periodic smooth function g of two variables can be approximated arbitrarily well by functions of the form

$$g_{n_{\text{tor}},m_{\text{pol}}}(\varphi,\theta) = \sum_{i=1}^{2n_{\text{tor}}+1} \sum_{j=1}^{2m_{\text{pol}}+1} a_{i,j} v_i(\varphi) w_j(\theta), v_i \in V^{n_{\text{tor}}}, w_j \in V^{m_{\text{pol}}}.$$

Hence, any surface $\boldsymbol{\Sigma}(\varphi,\theta)$ that can be smoothly parametrized in terms of periodic angles φ and θ can be approximated by

$$\boldsymbol{\Sigma}(\varphi,\theta) = \begin{bmatrix} \sum_{i,j} x_{i,j} v_i(\varphi) w_j(\theta) \\ \sum_{i,j} y_{i,j} v_i(\varphi) w_j(\theta) \\ \sum_{i,j} z_{i,j} v_i(\varphi) w_j(\theta) \end{bmatrix},$$

where the bounds on i and j in the double summation have been neglected for brevity. To describe stellarator relevant surfaces efficiently, we want to be able to exploit rotational symmetry. Denoting the rotation matrix by

$$M(\varphi) = \begin{bmatrix} \cos(2\pi\varphi) & -\sin(2\pi\varphi) & 0 \\ \sin(2\pi\varphi) & \sin(2\pi\varphi) & 0 \\ 0 & 0 & 1 \end{bmatrix},$$

a surface $\boldsymbol{\Sigma}$ satisfies n_{fp}-fold rotational symmetry if

$$\boldsymbol{\Sigma}(\varphi + 1/n_{\mathrm{fp}},\theta) = M(1/n_{\mathrm{fp}})\boldsymbol{\Sigma}(\varphi,\theta).$$

Hence, denoting

$$\tilde{\boldsymbol{\Sigma}}(\varphi,\theta) = M(-\varphi)\boldsymbol{\Sigma}(\varphi,\theta),$$

we note that $\tilde{\boldsymbol{\Sigma}}$ is $1/n_{\mathrm{fp}}$ periodic in φ, since

$$\tilde{\boldsymbol{\Sigma}}(\varphi + 1/n_{\mathrm{fp}},\theta) = M(-\varphi - 1/n_{\mathrm{fp}})\boldsymbol{\Sigma}(\varphi + 1/n_{\mathrm{fp}},\theta) = M(-\varphi)\boldsymbol{\Sigma}(\varphi,\theta) = \tilde{\boldsymbol{\Sigma}}(\varphi,\theta).$$

We conclude that a surface with n_{fp}-fold rotational symmetry can be approximated by functions of the form

$$\boldsymbol{\Sigma}(\varphi,\theta) = M(\varphi) \begin{bmatrix} \sum_{i,j} x_{i,j} v_i(n_{\mathrm{fp}}\varphi) w_j(\theta) \\ \sum_{i,j} y_{i,j} v_i(n_{\mathrm{fp}}\varphi) w_j(\theta) \\ \sum_{i,j} z_{i,j} v_i(n_{\mathrm{fp}}\varphi) w_j(\theta) \end{bmatrix}$$

B. Derivation of the major radius formula

The formula for the major radius is computed from the mean cross sectional area $\bar{A}$ and volume enclosed by the surface

$$R_{\mathrm{major}}(\mathbf{s}) = V(\mathbf{s})/(2\pi\bar{A}(\mathbf{s})),$$

where

$$\bar{A}(\mathbf{s}) = \frac{1}{2\pi}\int_{-\pi}^{\pi} A_\phi \; d\phi,$$

where A_ϕ is the surface's cross sectional area computed at the cylindrical angle ϕ. Note that there are multiple ways to evaluate this term. One possibility is to determine the cross sectional area at a fixed number of cylindrical angles ϕ, and then average. This is difficult with our surface representation, as it requires solving several nonlinear equations. We follow an alternative approach that is less complex implement and more elegant. The cross sectional area can be written as

$$A_\phi = \int_0^1 \tilde{R}\frac{\partial \tilde{Z}}{\partial \theta} \; d\theta,$$

where $\tilde{R}(\theta), \tilde{Z}(\theta)$ are points on the boundary of the cross section at the angle ϕ. Next, writing the above integrand in terms of the surface's Boozer angles, we have

$$A_\phi = \int_0^1 R(\varphi(\phi,\theta),\theta)\frac{\partial}{\partial\theta}[Z(\varphi(\phi,\theta),\theta)] \; d\theta,$$

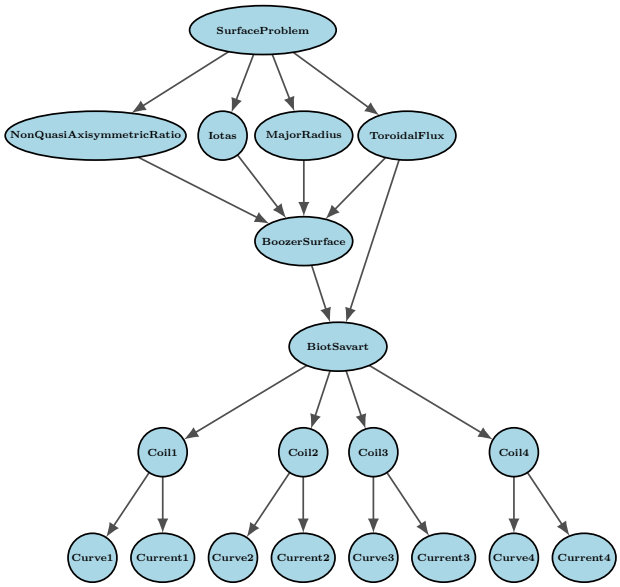

FIGURE 10. A simplified dependency graph of the various quantities used to define the optimization problem. In this model problem, there is one surface ("BoozerSurface" node) on which we are trying to optimize for quasi-symmetry ("NonQuasiAxisymmetricRatio" node), target a specific rotational transform ("Iotas" node), major radius ("MajorRadius" node) and toroidal flux ("ToroidalFlux" node). These penalty terms depend on on the magnetic field computed via the Biot Savart law ("BiotSavart" node) and on the geometry of the Boozer surface ("BoozerSurface" node). The Biot Savart node then depends directly on the four stellarator coil geometries and associated currents.

where $\tilde{R} = R(\varphi(\phi,\theta),\theta)$ and $\tilde{Z} = Z(\varphi(\phi,\theta),\theta)$. The expression for the average cross sectional area then becomes

$$\bar{A} = \frac{1}{2\pi}\int_{-\pi}^{\pi}\int_{0}^{1} R(\varphi(\phi,\theta),\theta)\frac{\partial}{\partial\theta}[Z(\varphi(\phi,\theta),\theta)]\ d\theta\ d\phi,$$

Instead of integrating over cylindrical ϕ, we complete the change of variables

$$(\phi,\theta) = (\mathrm{atan2}(y(\varphi,\theta),x(\varphi,\theta)),\theta).$$

After the change of variables, the average cross sectional area is

$$\bar{A} = \frac{1}{2\pi}\int_{0}^{1}\int_{0}^{1} R(\varphi,\theta)\left[\frac{\partial Z}{\partial\varphi}\frac{\partial\varphi}{\partial\theta} + \frac{\partial Z}{\partial\theta}\right]\ \det J\ d\varphi\ d\theta,$$

where $\det J$ is the determinant of the mapping's Jacobian.

C. Optimizable graph for efficient representation of the objective

The different penalty terms of the optimization problem in section 4.1 are implemented using the "optimizable" framework in SIMSOPT (Landreman *et al.* 2021). The dependencies between the various quantities in optimization problems result in a directed acyclical graph, which can be used for efficient caching, i.e., terms are only recomputed when quantities that they depend on are changed. The dependency graph for a simplified version of the optimization problem solved in section 4.1 is provided in Figure 10. A quantity associated to any of the nodes in the graph is only recomputed if any of its

ancestors have been modified. The full optimization problem solved in section 4.1 results in a much more complex dependency graph that we do not show here for brevity.