

PAPER • OPEN ACCESS

Global stellarator coil optimization with quadratic constraints and objectives

To cite this article: Lanke Fu *et al* 2025 *Nucl. Fusion* **65** 026045

View the [article online](#) for updates and enhancements.

You may also like

- [The Quadruple-lined, Doubly Eclipsing System V482 Persei](#)
Guillermo Torres, Claud H. Sandberg Lacy, Francis C. Fekel *et al.*
- [TIC 454140642: A Compact, Coplanar, Quadruple-lined Quadruple Star System Consisting of Two Eclipsing Binaries](#)
Veselin B. Kostov, Brian P. Powell, Guillermo Torres *et al.*
- [The TWA 3 Young Triple System: Orbits, Disks, Evolution](#)
Kendra Kellogg, L. Prato, Guillermo Torres *et al.*

Global stellarator coil optimization with quadratic constraints and objectives

Lanke Fu^{1,2,*} , Elizabeth J. Paul³ , Alan A. Kaptanoglu⁴ and Amitava Bhattacharjee²

¹ Princeton Plasma Physics Laboratory, Princeton, NJ, United States of America

² Department of Astrophysical Sciences, Princeton University, Princeton, NJ, United States of America

³ Department of Applied Physics, Columbia University, New York, NY, United States of America

⁴ Courant Institute of Mathematical Sciences, New York University, New York, NY, United States of America

E-mail: ffu@pppl.gov

Received 15 August 2024, revised 2 December 2024

Accepted for publication 9 January 2025

Published 22 January 2025



CrossMark

Abstract

Most present stellarator designs are produced by costly two-stage optimization: the first for an optimized equilibrium, and the second for a coil design reproducing its magnetic configuration. Few proxies for coil complexity and forces exist at the equilibrium stage. Rapid initial state finding for both stages is a topic of active research. Most present convex coil optimization codes use the least square winding surface method by Merkel (NESCOIL), with recent improvements in conditioning, regularization, sparsity, and physics objectives. While elegant, the method is limited to modeling the norms of linear functions in coil current. We present QUADCOIL, a global coil optimization method that targets combinations of linear and quadratic functions of the current. It can directly constrain and/or minimize a wide range of physics objectives unavailable in NESCOIL and REGCOIL, including the Lorentz force, magnetic energy, curvature, field-current alignment, and the maximum density of a dipole array. QUADCOIL requires no initial guess and runs nearly $10^2 \times$ faster than filament optimization. Integrating it in the equilibrium optimization stage can potentially exclude equilibria with difficult-to-design coils, without significantly increasing the computation time per iteration. QUADCOIL finds the exact, global minimum in a large parameter space when possible, and otherwise finds a well-performing approximate global minimum. It supports most regularization techniques developed for NESCOIL and REGCOIL. We demonstrate QUADCOIL's effectiveness in coil topology control, minimizing non-convex penalties, and predicting filament coil complexity with three numerical examples.

Keywords: stellarator, coils, winding surface, optimization, convex relaxation

(Some figures may appear in colour only in the online journal)

* Author to whom any correspondence should be addressed.



Original Content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

1. Introduction

Stellarators are attractive three-dimensional (3D) fusion devices that promise steady-state, disruption-free operation. Unlike tokamaks, stellarators generate rotational transform by external coils rather than a plasma current. The coil system of a stellarator reactor determines its cost and performance dominantly [1].

Traditionally, stellarator designs are produced by two-stage optimization: the first stage for an equilibrium with MHD stability and good particle confinement, and the second stage for a coil set that accurately reproduces the equilibrium magnetic field. Both stages are high-dimensional in parameter space, non-linear, and ill-posed optimization problems. The coil stage, specifically, is an inverse Biot–Savart problem that does not have a unique solution. For a vacuum configuration, a coil optimization code typically attempts to generate a current configuration, so that B_{normal} zero at the plasma boundary. For a finite beta equilibrium, the code will target a non-zero B_{normal} distribution at the plasma boundary calculated using the virtual casing principle.

Presently, there are two main methods for coil optimization: filament and winding surface. The filament method treats coils as discrete space curves. Examples of filament codes include COILOPT [2], FOCUS [3], SIMSOPT [4], and DESC [5]. This method has been shown to generate realistic coil sets. However, because it formulates coil optimization as a non-convex optimization problem, the filament method suffers from many common issues in high-dimensional, multi-objective non-convex optimization. Non-convex optimization is generally NP-hard. Good convergence typically requires careful manual tuning on the initial state and penalty weights. Moreover, in practice, several filament subproblems are solved sequentially, starting with low-resolution solutions that are used to initialize optimization at increasing resolutions. The coil number and topology in the filament method are also fixed during the optimization. Due to these disadvantages, designing a coil set using the filament method is often computationally costly and labor-intensive.

The winding surface method approximates a coil set as a sheet current \mathbf{K}' on a prescribed winding surface, then cuts the sheet current into discrete coils by choosing a set of current paths and giving each path an appropriate current. Examples of such codes include NESCOIL [6], ONSET [7], and REGCOIL [8]. Compared to the filament method, this method is less realistic, but it formulates coil optimization as a convex optimization problem. Therefore, the winding surface method has a theoretical complexity of $\mathcal{O}(n_{\text{sample}}n_{\text{dof}}^2)$, where n_{sample} is the number of sample points in the field error calculation, and n_{dof} is the number of degrees of freedom in \mathbf{K}' . Although the method requires an arbitrary choice of winding surface, it always converges to the global optimum regardless of the surface choice. In practice, the filament method and winding surface method are often used in complement. The filament method produces the final coil set, whereas the winding surface method provides initial states for the filament method [9].

The winding surface method is naturally suitable as a proxy for coil complexity in the equilibrium stage. The method has been used to study the impact of plasma parameters on coil complexity [10]. However, despite recent improvements [8, 11–13], its formulation greatly limits the choice of objectives and constraints. All existing winding surface methods are limited to minimizing the L-1 or L-2 norms of linear functions of \mathbf{K}' . This greatly limits the choice of physics and engineering metrics. All existing winding surface methods are limited to performing unconstrained optimization. This limits its control over the topology of \mathbf{K}' . As a result, the sheet current may contain complex features that make coil-cutting difficult.

We present QUADCOIL, a reformulation of the winding surface method. QUADCOIL is based on the convex relaxation of quadratic programs. It enables global coil optimization using a non-convex quadratic objective under non-convex quadratic constraints, with a similar speed to existing winding surface methods. QUADCOIL can directly target a wide range of new metrics. It offers greater control over the topology of \mathbf{K}' , allowing automated coil cutting without human intervention. Unlike existing winding surface methods, QUADCOIL requires no initial guess and minimal parameter tuning. This makes it suitable for initial state generation or as a coil complexity metric in the equilibrium optimization loop. Although it is tuned for fusion applications, QUADCOIL is fundamentally an inverse Biot–Savart code and can be used for other magnetostatic optimization problems, such as accelerator magnet design.

Notably, in recent years, single-stage stellarator optimization has risen in popularity. Single-stage optimization simultaneously optimizes the stellarator equilibrium and coil set to find a compromise between physics and coil complexity requirements [14–17]. Alternatively, the approach can also design stellarators that are robust to coil manufacturing errors. While the single-stage approach has successfully produced a large number of optimized equilibria with quasisymmetry (QS), its parameter space may be needlessly large. Because one equilibrium can be reproduced by many different coil configurations, similar equilibria can potentially appear in multiple, distant regions of the large parameter space. This may increase the number of local minima compared to those in either stage of the two-stage approach, which makes the optimization process more costly and harder to converge. In comparison, QUADCOIL connects the coil and equilibrium stage during a two-stage optimization as a coil complexity metric. Like existing coil optimization codes, QUADCOIL is agnostic to the type of force balance inside the equilibrium. We believe our approach can be used to identify vacuum or finite-beta equilibria that balance plasma performance and coil complexity, similar to a single-stage optimization, but without introducing additional degrees of freedom.

Our paper is organized as follows. Section 2.1 introduces the theory behind existing winding surface methods and surveys recent developments. Section 2.2 discusses QUADCOIL's formulation and motivations. Sections 3.1–3.3 demonstrate QUADCOIL's effectiveness in coil topology

control, minimizing non-convex penalty, and predicting filament coil complexity with three numerical examples.

2. Theory

2.1. Winding surface method

In this section, we will briefly introduce the formulation shared by existing winding surface methods, their advantages, and their limitations.

Stellarator coil optimization, at its core, is an inverse Biot–Savart problem. Typically, this problem is non-convex, due to the $|\mathbf{r} - \mathbf{r}'|^{-3}$ non-linearity in the Biot–Savart Law (In this paper, primed letters will denote coil-related quantities, and unprimed letters will denote plasma-related quantities)

$$\mathbf{B}(\mathbf{r}) = \int d^2a' \frac{\mathbf{I}' \times (\mathbf{r} - \mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|^3}. \quad (1)$$

Here, \mathbf{I}' is the coil current vector, and \mathbf{r} and \mathbf{r}' are locations at the plasma boundary and the coils. The winding surface method circumvents this non-linearity by fixing the location of the current. Because \mathbf{K}' is divergence-free, we can represent it with the gradient of a current potential Φ' [6]:

$$\nabla \cdot \mathbf{K}' = 0 \Leftrightarrow \mathbf{K}' = \hat{\mathbf{n}}' \times \nabla \Phi', \quad (2)$$

$$\begin{aligned} \mathbf{B}(\mathbf{r}) &= \frac{\mu_0}{4\pi} \int d^2a' \frac{\mathbf{K}' \times (\mathbf{r} - \mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|^3} \\ &= \frac{\mu_0}{4\pi} \sum_{\ell'=1}^{n_p} \int_0^{2\pi} d\theta' \int_0^{2\pi/n_p} d\zeta' \frac{1}{|\mathbf{r} - \mathbf{r}'|^3} \left(\frac{\partial \mathbf{r}'}{\partial \theta'} \frac{\partial \Phi'}{\partial \zeta'} \right. \\ &\quad \left. - \frac{\partial \mathbf{r}'}{\partial \zeta'} \frac{\partial \Phi'}{\partial \theta'} \right) \times (\mathbf{r} - \mathbf{r}'). \end{aligned} \quad (3)$$

where $\hat{\mathbf{n}}'$ is the unit normal of the winding surface. Generally, the current potential Φ' consists of three components:

$$\Phi' = \Phi'_{sv}(\zeta', \theta') + \frac{G\zeta'}{2\pi} + \frac{I\theta'}{2\pi}. \quad (4)$$

Here, ζ', θ' are arbitrary toroidal and poloidal angles on the winding surface and Φ'_{sv} is the single-valued component of Φ' , the degrees of freedom we optimize. G, I are the net toroidal and poloidal currents. The value of I is determined by the toroidal flux of the equilibrium, whereas the value G is a free parameter that determines whether \mathbf{K}' forms primarily poloidal or helical coils. For brevity, we will abbreviate Φ'_{sv} as Φ' in the rest of this paper.

To visualize Φ' , the current \mathbf{K}' always flows along contours of equal Φ' . The ratio between I and G determines the helicity of the net current. Alternatively, the current potential Φ' has the unit of surface magnetic dipole moment density, and its single-valued component, Φ'_{sv} , can be considered as a dipole sheet oriented perpendicular to the winding surface. Therefore, the winding surface approach is suitable for modeling a variety of coil types, including poloidal coils, helical coils, windowpane coils, and dipole arrays. Using multiple

winding surfaces with different objectives, constraints, and net currents, it may also be possible to model a combination of different coil types.

Since the Biot–Savart Law is linear in \mathbf{K}' , we can minimize f_B , the squared flux at the plasma boundary, as a linear least square problem,

$$\min_{\Phi'} f_B(\Phi') = \min_{\Phi'} \|A_B \Phi' + b_B\|_2, \quad (5)$$

$$f_B \equiv \int_{\text{plasma}} d^2a |\mathbf{B}_{\text{coil}} \cdot \hat{\mathbf{n}} - B_T|^2. \quad (6)$$

Here, $\hat{\mathbf{n}}$ is the surface normal of the plasma surface, \mathbf{B}_{coil} is the magnetic field produced by the stellarator coils, and B_T is the target normal field at the plasma boundary. This is a well-understood convex optimization problem with a closed-form solution.

The optimization problem in (5) is the basis of NESCOIL, which played important roles in the coil design of W7X and NCSX. However, the ill-posed nature of the problem means that the induction matrix A_B is ill-conditioned. Without proper regularization, optima of (5) typically contain unrealistically high peak current, as well as areas with dense, smoothly varying current that poorly approximate discrete coils.

Recent improvements in winding surface methods include better conditioning, regularization, sparsity promotion, linear physics objectives, and finite element bases. Boozer [11] performed truncated singular value decomposition (TSVD) on A_B to remove small-scale, high-amplitude current modes with little impact on f_B . In the code REGCOIL, Landreman [8] introduces a Tikhonov regularization term,

$$\begin{aligned} &\min_{\Phi'} \left[f_B(\Phi') + \lambda_2 \int d^2a \|\mathbf{K}'\|_2 \right] \\ &= \min_{\Phi'} (\|A_B \Phi' + b_B\|_2 + \lambda_2 \|A_K \Phi' + b_K\|_2), \end{aligned} \quad (7)$$

regularization weight $\lambda_2 \geq 0$.

This modifies (5) into a ridge regression problem and effectively reduces the complexity of the resulting current. Elder [18] introduces L-1 regularization,

$$\begin{aligned} &\min_{\Phi'} \left[f_B(\Phi') + \lambda_1 \int d^2a \|\mathbf{K}'\|_1 \right] \\ &= \min_{\Phi'} (\|A_B \Phi' + b_B\|_2 + \lambda_1 \|A_K \Phi' + b_K\|_1), \end{aligned} \quad (8)$$

regularization weight $\lambda_1 \geq 0$.

The new regularization term promotes sparsity in \mathbf{K}' and modifies (5) into a lasso regression problem. Both types of regularization are convex and can be applied to other linear functions of \mathbf{K}' . Kaptanoglu *et al* [13] introduce a non-convex, L-0 regularized current voxel method for coil topology optimization. However, its non-smooth nature places it outside the scope of this work.

These improvements highlight the two remaining shortcomings of existing winding surface methods. The first limitation is in the choice of objective. Both (7) and (8) are still

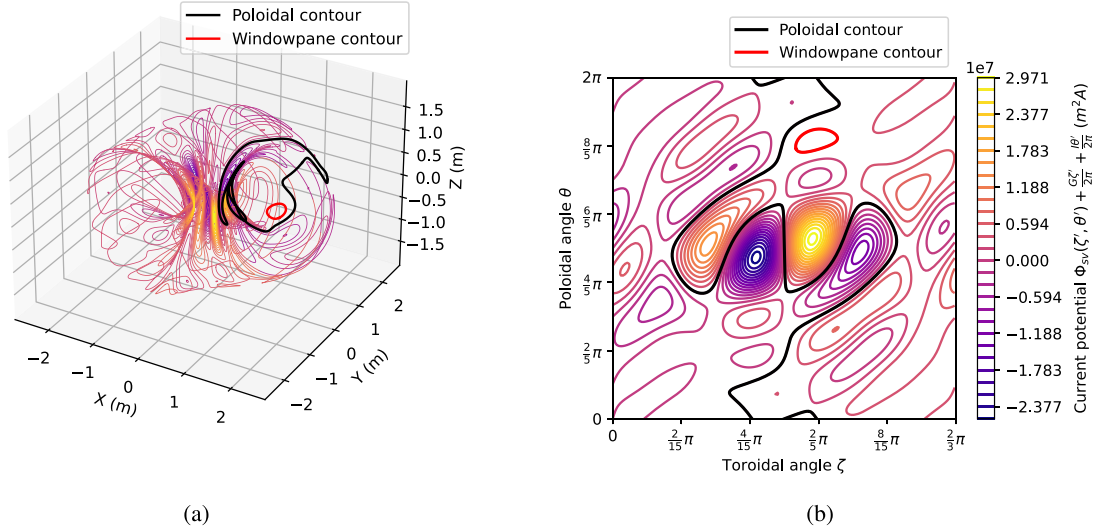


Figure 1. Φ' contours in (X, Y, Z) (left) and (ζ, θ) (right) of the NESCOIL optimum for NCSX, using a winding surface ≈ 2 minor radii from the plasma generated using the procedure in section 2.4, and 4 Fourier harmonics each in the poloidal and toroidal directions. Φ' has the units of magnetic dipole moment density. Note the presence of both poloidal (links with the plasma) and windowpane (does not link with the plasma) contours.

limited to minimizing the L-2 or L-1 norm of linear functions of \mathbf{K}' . The second limitation is that all existing winding surface methods are unconstrained. This makes it difficult to control the topology of the resulting surface current, beyond specifying the net poloidal current G and the net toroidal current I . Specifically, NESCOIL and REGCOIL solutions usually contain a mix of poloidal, helical, and windowpane Φ' contours. This is especially likely on a compact device, as figure 1 shows. Because windowpane contours can be arbitrarily small, they can greatly complicate the coil-cutting process. To avoid unrealistic small windowpane coils, one often needs to trace and classify a large number of contours, or manually pick a set of realistic seeming coils.

2.2. QUADCOIL

Our algorithm, QUADCOIL, aims to address both drawbacks by allowing global optimization of a surface current by minimizing a non-convex quadratic objective under non-convex quadratic constraints. In this section, we discuss its formulation and characteristics.

QUADCOIL formulates the winding surface problem as a quadratically constrained quadratic program (QCQP):

$$\begin{aligned} & \min_{\Phi', a, b} [f(\Phi', a) + c^T b] \\ & \text{subject to} \\ & g_j(\Phi', a) + d_j^T b \leq e_j, \\ & h_k(\Phi', a) + m_k^T b = n_k, \\ & \text{for all } j = 1, \dots, n_{\text{ineq}}, k = 1, \dots, n_{\text{eq}}. \end{aligned} \quad (9)$$

Here, f , g_j are quadratic functions of Φ' and a . n_{ineq} and n_{eq} are the total number of inequality and equality constraints. $\Phi' \in \mathbb{R}^{n_{\Phi'}}$ is the discretized current potential. $a \in \mathbb{R}^{n_a}$, $b \in \mathbb{R}^{n_b}$

are auxiliary variables for constructing complex problems. c , d_j and e_j are constants. Unlike in NESCOIL or REGCOIL, f or g_j may be convex or non-convex. This allows QUADCOIL to target a greater range of objectives and constraints.

By allowing non-convex quadratic objectives and constraints, QUADCOIL can target important engineering and physics objectives unavailable to existing winding surface methods. These include non-convex, quadratic terms such as stored magnetic energy,

$$E = \frac{1}{2\mu_0} \int_{\infty} d^3\mathbf{r} |\mathbf{B}|^2, \quad (10)$$

self-Lorentz force [19],

$$\begin{aligned} \mathbf{L}(\mathbf{y}) = & -\frac{\mu_0}{4\pi} \oint_S d^2\mathbf{x} \frac{1}{|\mathbf{y}-\mathbf{x}|} \{ \nabla_{\mathbf{x}} \cdot [\pi_{\mathbf{x}} \mathbf{K}'(\mathbf{y})] + \pi_{\mathbf{x}} \mathbf{K}'(\mathbf{y}) \cdot \nabla_{\mathbf{x}} \} \mathbf{K}'(\mathbf{x}) \\ & + \frac{\mu_0}{4\pi} \oint_S d^2\mathbf{x} [\mathbf{K}'(\mathbf{y}) \cdot \mathbf{n}(\mathbf{x})] \frac{(\mathbf{y}-\mathbf{x}) \cdot \mathbf{n}(\mathbf{x})}{|\mathbf{y}-\mathbf{x}|^3} \mathbf{K}'(\mathbf{x}) \\ & + \frac{\mu_0}{4\pi} \oint_S d^2\mathbf{x} \frac{1}{|\mathbf{y}-\mathbf{x}|} \{ \mathbf{K}'(\mathbf{y}) \cdot \mathbf{K}'(\mathbf{x}) \nabla_{\mathbf{x}} \cdot \pi_{\mathbf{x}} + \nabla_{\mathbf{x}} [\mathbf{K}'(\mathbf{y}) \cdot \mathbf{K}'(\mathbf{x})] \} \\ & - \frac{\mu_0}{4\pi} \oint_S d^2\mathbf{x} [\mathbf{K}'(\mathbf{y}) \cdot \mathbf{K}'(\mathbf{x})] \frac{(\mathbf{y}-\mathbf{x}) \cdot \mathbf{n}(\mathbf{x})}{|\mathbf{y}-\mathbf{x}|^3} \mathbf{n}(\mathbf{x}), \\ \pi_{\mathbf{x}} \equiv & \mathbb{I} - \mathbf{n}(\mathbf{x}) \mathbf{n}(\mathbf{x}), \end{aligned} \quad (11)$$

the curvature proxy $\mathbf{K}' \cdot \nabla \mathbf{K}'$, and coil-field alignment $\mathbf{K}' \cdot \mathbf{B}$, potentially useful for estimating the critical current of high-temperature superconductor (HTS) coils. By combining constraints and auxiliary variables, QUADCOIL can also target functional forms unavailable to existing methods. Examples include the maximum density of a dipole array, $\max_{\zeta', \theta'} |\Phi'_{\text{sv}}|$, and the local magnetic field error, $\max_{\zeta', \theta'} |\mathbf{B}|^2$.

In this paper, we will focus on the quantity,

$$f_{\kappa}^{\infty} \equiv \max_{\zeta', \theta'} \| (\mathbf{K}' \cdot \nabla \mathbf{K}')_{(r, \Phi', z)} \|_{\infty}, \quad (12)$$

as the primary example. This quantity is a quadratic proxy for the maximum curvature of the surface current,

$$\kappa_{\max} = \max_{\zeta', \theta'} \frac{\|\mathbf{K}' \cdot \nabla \mathbf{K}' \times \mathbf{K}'\|_2}{\|\mathbf{K}'\|_2^3}, \quad (13)$$

as well as the curvature of a filament coil with current vector \mathbf{I} ,

$$\kappa = \frac{\|\mathbf{I} \cdot \nabla \mathbf{I}\|_2}{I^2}. \quad (14)$$

The construction of f_{κ}^{∞} will be discussed in greater detail in appendix A.

With proper choice of f, g_j, a , and b , (9) can represent a wide range of coil optimization problems. These include the REGCOIL problem (7), the L-1-regularized problem (8), and interesting new problems, such as minimizing the stored magnetic energy $E(\Phi')$ under a local error constraint:

$$\begin{aligned} & \min_{\Phi'} E(\Phi') \\ & \text{subject to} \\ & |\mathbf{B} \cdot \hat{\mathbf{n}}| \leq B_{\max}. \end{aligned} \quad (15)$$

While a general QCQP is NP-hard, the exact global optimum of a ‘nearly convex’ QCQP can usually be found via its relaxation as a conic programming problem. This procedure is called Shor relaxation [20], and is widely used for QCQPs outside stellarator optimization. Shor relaxation is uniquely suited for stellarator coil optimization because the primary consideration, field error, is often formulated as a convex quadratic term (i.e. f_B). This allows QUADCOIL to quickly find the global optimum of (9) and avoid the time-consuming parameter tuning associated with constrained, non-convex optimization.

QUADCOIL combines Shor relaxation and a non-convex optimizer to perform rapid global optimization of stellarator coils. We discuss the properties of Shor relaxation, the implementation of QUADCOIL, the exactness of QUADCOIL solutions, as well as its limitations, in section 2.3. The compatibility between Shor relaxation and the coil optimization problem will be further explored in section 3.3.

QUADCOIL applies to all linear discretizations of Φ' . For simplicity, this paper uses the same Fourier representation as in REGCOIL and NESCOIL:

$$\Phi'(\theta', \zeta') = \sum_j \Phi'_j \begin{pmatrix} \sin \\ \cos \end{pmatrix}_j (m_j \theta' - n_j \zeta'), \quad (16)$$

where j is a uniform index for the poloidal, toroidal mode numbers (m, n) , and the choice of sin and cos. Similar to NESCOIL and REGCOIL, QUADCOIL supports all types of coil topology.

2.3. Shor relaxation

Shor relaxation rewrites a QCQP as a conic program (CP). This is a well-studied class of convex optimization problems

and can be solved to arbitrary precision in polynomial time using interior point methods. A convex QCQP always has an exactly equivalent CP. While a non-convex QCQP does not have an exactly equivalent CP, it often has a CP that shares its global optimum or gives a good approximation to its global optimum. The detailed procedure for relaxing (9) is summarized in appendix C.

There is a rich literature on the classification of QCQP with exact Shor relaxation, beyond the scope of this paper. Empirically, Shor relaxation typically works best when a problem is nearly convex. This is usually the case in winding surface optimization because the dominant factor in coil optimization is almost always the accuracy of the magnetic field, a convex quantity. Notably, besides f_B , many other forms of magnetic field objectives can be written as convex QCQPs. One example is the maximum local field error:

$$f \text{ or } g = \max_{\zeta', \theta'} \|\mathbf{B} \cdot \mathbf{n}\|_2 \Leftrightarrow \quad (17)$$

$$f \text{ or } g = a, \text{ subject to} \quad (18)$$

$$(\|\mathbf{B} \cdot \mathbf{n}\|_2)_{\zeta', \theta'} \leq a, \quad (19)$$

$$a \geq 0. \quad (20)$$

While it is tricky to predict from f, g_j, c_j and d_j whether (9) has an exact Shor relaxation, once the solution to the CP is obtained, it is easy to numerically test its exactness, requiring only the SVD of a $(n_{\Phi'} + n_a + 1) \times (n_{\Phi'} + n_a + 1)$ matrix. In practice, the CP of a typical QUADCOIL problem (40 Φ' harmonics, 1024 quadrature points per field period) takes only core seconds to solve. Therefore, it is practical to solve the CP first and then test its exactness. We will discuss the mechanism of the optimality test in greater detail in appendix C. In section 3.3, we will validate its effectiveness, and demonstrate the compatibility between Shor relaxation and stellarator coil optimization as part of a parameter space study.

Using these properties, QUADCOIL can substantially speed up constrained, non-convex winding surface problems. The algorithm works as follows:

1. Upon receiving a QCQP coil optimization problem, QUADCOIL first solves its Shor relaxation and tests its exactness.
2. If the solution is exact, QUADCOIL returns the answer.
3. If the solution is not exact, QUADCOIL initializes a local optimization solver, using the relaxation as the initial guess.

The implementation of QUADCOIL in this paper uses CVXPY [21, 22] and MOSEK [23] as the CP solver. Alternatively, when no constraint is present, (9) can also be efficiently solved by local optimizers. The numerical example in section 3.2 is one such unconstrained variation of (9). In section 3.2, we will compare Shor relaxation to two popular unconstrained local optimization algorithms, ADAM [24], a stochastic gradient descent algorithm, and L-BFGS-B [25], a quasi-Newton method. We will also briefly discuss the comparative advantage of each solver for different variations of (9).

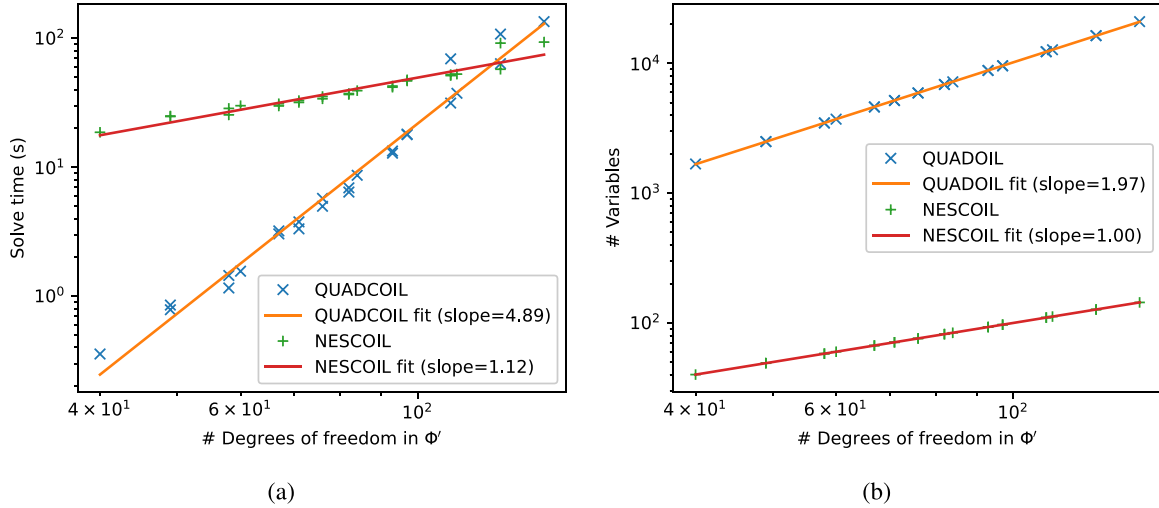


Figure 2. Time t required to solve (21) (left) and the number of free variables (right) for NCSX as a function of $n_{\Phi'}$. Note the power-of-two difference in the numbers of free variables in QUADCOIL and NESCOIL.

For simplicity, in this paper, we have only implemented step 3 for unconstrained problems using L-BFGS-B. The choice and tuning of a constrained local optimization solver for this step requires further work. However, section 3.3 shows that even without step 3, approximate solutions from QUADCOIL can still predict filament coil complexity.

Shor relaxation comes with one drawback. The number of free variables in the relaxed CP scales as $\mathcal{O}[(n_{\Phi'} + n_a)^2 + n_b]$. As a result, compared to existing current potential methods, QUADCOIL scales poorly with high $n_{\Phi'}$ and n_a . This may become a problem in future works applying finite element basis to QUADCOIL. See figure 2 for the solve times and variable numbers of (24) in QUADCOIL as a function of n . ($m = 0$ in this problem).

2.4. Generating winding surfaces

Procedures for generating a well-behaved winding surface have received limited attention in the current literature. Existing winding surface codes typically generate the winding surface by uniformly offsetting the plasma surface, or its poloidal cross sections, along the normal direction [8, 10]. For a highly shaped plasma surface, this method can produce a winding surface with self-intersections, regions with very large principal curvatures, and unevenly spaced quadrature points. Therefore, existing works often iterate the winding surface shape over many repeated NESCOIL/REGCOIL solves to create a sufficiently well-behaved solution [7, 9, 12]. Fundamentally, the iteration process is a non-convex optimization problem, similar to filament coil optimization. Although it has less degrees of freedom than filament optimization, one can argue that it negates the winding surface method's advantages in speed and solution uniqueness.

An ill-behaved winding surface is especially detrimental to the minimization of f_{κ}^{∞} . Speed also is imperative for QUADCOIL to become a valid coil complexity proxy. Under these motivations, we present a new, non-iterative winding

surface generator for QUADCOIL. Our method can directly produce a well-behaved winding surface without repeated QUADCOIL solves. The method follows these five steps:

1. Offset quadrature points on the plasma surface along the unit normal by d_{cs} .
2. Evaluate the cross sections of the surface at uniform intervals in the Cartesian toroidal angle, ζ' .
3. Find the convex hull of the cross-section.
4. Fit a periodic cubic spline to the convex hull and redefine quadrature points at equal arc lengths.
5. Fit a surface to the quadrature points.

The poloidal coordinate on the winding surface, θ' , is proportional to the arc length of the cubic spline. We choose the zeros of θ' as the outboard point with Z closest to the splines' centers of weight. Figure 3 compares our method with the uniform offset method. Note that our method creates a smoother surface with no self-intersections and uniformly spaced quadrature points.

This method for generating winding surfaces has one potential drawback. Empirically, the field error in a stellarator is the most sensitive to the coil-plasma distance on the concave regions of the plasma surface. Using the convex hull of cross sections, therefore, can potentially increase the optimal field error. Yet, one may argue that a physical vessel with no concave region is simpler to engineer. The method for producing the optimal winding surface is beyond the scope of this paper. As section 3.3 shows, our method is sufficiently well-behaved across a variety of equilibria and optimization targets for QUADCOIL to serve as a coil curvature proxy.

3. Numerical examples

In this section, we demonstrate QUADCOIL's potential as an initial state tool and coil complexity proxy with three numerical examples. In section 3.1, we demonstrate an inequality

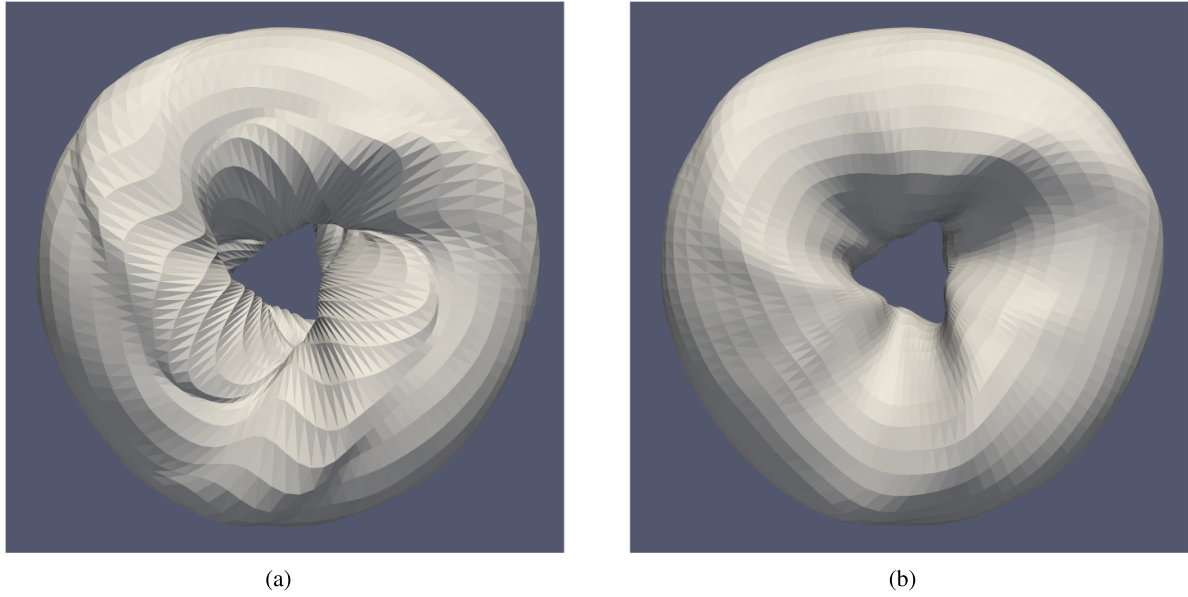


Figure 3. NCSX winding surface generated by uniformly offsetting quadrature points on the plasma surface along the normal direction (left) vs. surface generated by QUADCOIL (right).

constraint that gives the option to eliminate all windowpane Φ' contours and substantially simplify the coil-cutting process. In section 3.2, we show that QUADCOIL can outperform REGCOIL in coil simplification by directly targeting a non-convex penalty. In section 3.3, we validate QUADCOIL's effectiveness as a coil complexity proxy. We observe a correlation between the optimal normalized field error produced by QUADCOIL and filament optimizations across 4436 combinations of different equilibrium, maximum curvature, coil-coil spacing, and coil-plasma spacing [26].

3.1. Topology control

QUADCOIL gives the option to create a current configuration with purely poloidal Φ' contours. This is achieved by setting $G = 0$, and solving:

$$\min_{\Phi'} f_B(\Phi'), \quad (21)$$

$$\text{subject to } K'_\theta(\Phi') \text{sign}(I) \geq 0. \quad (22)$$

Here, $\text{sign}(I)$ is the sign of the net poloidal current. (22) is a necessary condition for all Φ' contours to be poloidal. It prevents the sheet current \mathbf{K}' from flowing against the direction of the net poloidal current, effectively forbidding windowpane currents.

A similar constraint can be constructed for helical coils:

$$\mathbf{K}'(\Phi') \cdot \hat{\mathbf{l}}_{\text{helical}} \geq 0, \quad (23)$$

where $\hat{\mathbf{l}}_{\text{helical}}$ is a unit vector tangent to the direction of the net helical current.

It is difficult to obtain a similar collection of poloidal current paths with REGCOIL because the REGCOIL penalty does not directly target the sign of poloidal current. The first

method is to evaluate a large number of contours during the coil-cutting process and choose a subset that has the desired topology and provides a sufficiently accurate magnetic field. This is computationally costly and often requires human input. Alternatively, one can search for a regularization parameter λ_2 in (7) that produces the lowest- f_B current configuration with the desired topology, although there is no guarantee a given topology will be achieved. Moreover, this requires a large number of REGCOIL solves and can converge slowly without a good initial guess for λ_2 .

Compared to REGCOIL, QUADCOIL can produce a purely poloidal current configuration with lower f_B in a fraction of the computation time. Figure 4 compares two such current configurations for NCSX [27], obtained with QUADCOIL and REGCOIL. The winding surface is located 2 minor radii away from the plasma. We obtain the configuration in figure 4(a) by solving (21) in QUADCOIL, without any additional regularization. This configuration has $f_B = 0.607 \text{ T}^2 \text{ m}^2$ and requires 1.77 s to solve. We obtain the configuration in figure 4(b) by performing a binary search in $\log_{10}(\lambda_2)$, to find the lowest- f_B REGCOIL solution satisfying $K'_\theta(\Phi') \text{sign}(I) \geq 0$. The search starts at $\log_{10}(\lambda_2) = -50$ and 1, and terminates when the relative change in f_B falls below 0.1%. The binary search requires 15 REGCOIL solves and 43.50 s to complete, substantially slower than QUADCOIL. The configuration has $f_B = 0.792 \text{ T}^2 \text{ m}^2$, 30.4% higher than the QUADCOIL value.

3.2. Non-convex penalty

In this section, we show that QUADCOIL can outperform REGCOIL when directly targeting a non-convex objective. In this section, we design helical coils for HSX, and modular coils for NCSX and W7X, by solving




Figure 4. The first row shows the Φ' contours of the lowest- f_B QUADCOIL and REGCOIL configurations with no windowpane current. The second row shows their $|B_{\text{normal}}|$ distribution. Both configurations use the same winding surface and number of Fourier harmonics as figure 1. Note that the REGCOIL configuration has higher field error, because Tikhnov regularization does not directly target the sign of the poloidal current, and therefore overconstrains.

$$\min_{\Phi'} [f_B(\Phi') + \lambda_\kappa f_\kappa^\infty(\Phi')]. \quad (24)$$

All three cases share the same winding surface located 2 minor radii from the plasma.

QUADCOIL consistently outperforms REGCOIL when targeting specific non-convex penalties. Figure 5 compares the trade-off between f_B and f_κ^∞ across REGCOIL and QUADCOIL using 3 different solvers:

1. A hybrid solver combining Shor relaxation and L-BFGS-B, as discussed in section 2.3.
2. Directly solving with L-BFGS-B, the quasi-Newton algorithm.
3. Directly solving with ADAM, the stochastic gradient descent algorithm.

As figure 5 shows, with all three solvers, QUADCOIL consistently produces lower f_κ^∞ than REGCOIL at equal f_B . Note that while a single QUADCOIL case runs in a similar time as a single REGCOIL case, producing a configuration with a target f_B requires a costly λ_2 search similar to the example in section 3.1. Therefore, the actual time required to obtain a desired current configuration using REGCOIL can be significantly longer than the per-case solve time reported in table 1.

We now compare the three different solvers for unconstrained QUADCOIL. Although all three solvers yield comparably accurate results, ADAM is 5–20 \times slower than the hybrid solver and L-BFGS-B. This suggests that the landscape of (24) may be sufficiently smooth and nearly convex that stochastic gradient descent is not necessary. As discussed in section 2.3, directly solving with L-BFGS-B is also likely to scale better than Shor relaxation to a higher number of degrees of freedom. Therefore, we believe that the L-BFGS-B solver

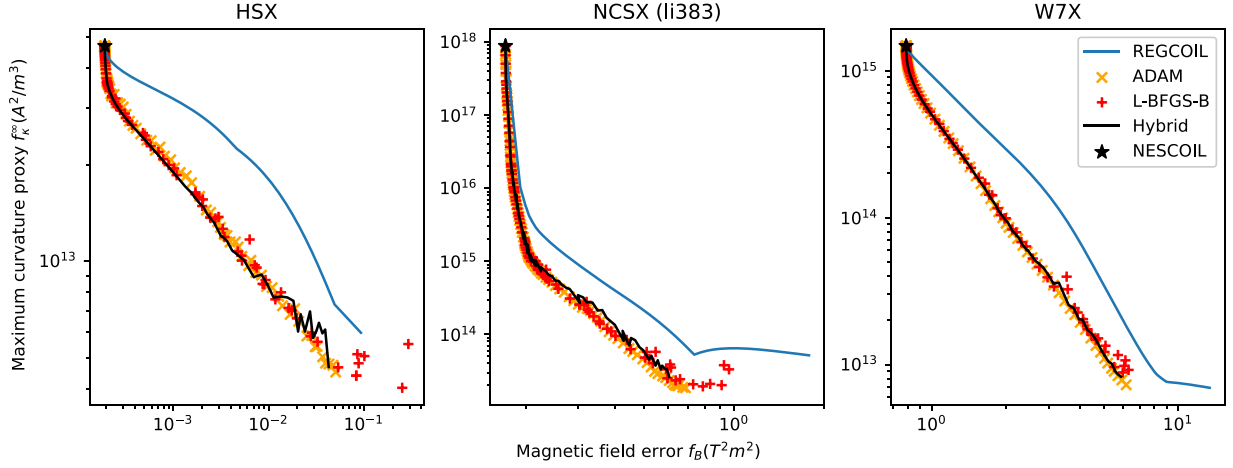


Figure 5. Trade-off plots between the squared flux f_B and the curvature proxy f_κ^∞ . This figure includes 150 REGCOIL runs, and 100 runs each for QUADCOIL using the hybrid, ADAM, and L-BFGS-B solvers. The REGCOIL runs sweep λ_2 exponentially. The QUADCOIL runs solve (24), sweeping λ_κ exponentially. For ADAM or L-BFGS-B, each marker represents a scan point. For visual clarity, the REGCOIL and QUADCOIL hybrid scan points are not marked.

Table 1. Minimum, maximum, and median solve time comparison of QUADCOIL, L-BFGS-B, and ADAM minimizing (24). Note that the REGCOIL solve times in this table are per case. To find a REGCOIL solution with a target f_B will require a λ_2 search, and will be substantially slower than the numbers in this table.

	REGCOIL	QUADCOIL (hybrid)	QUADCOIL (L-BFGS-B)	QUADCOIL (ADAM)
HSX	12.6–12.7 s, Med(t) = 12.7 s	0.4–15.8 s, Med(t) = 1.4 s	1.0–19.3 s, Med(t) = 3.4 s	5.0–32.7 s, Med(t) = 10.5 s
NCSX	9.5–10.9 s, Med(t) = 9.5 s	0.2–13.1 s, Med(t) = 4.7 s	2.3–12.5 s, Med(t) = 5.0 s	9.6–73.4 s, Med(t) = 25.8 s
W7X	16.0–18.7 s, Med(t) = 16.0 s	0.3–12.5 s, Med(t) = 1.5 s	1.2–16.9 s, Med(t) = 3.6 s	9.8–63.2 s, Med(t) = 34.9 s

is preferable at high resolutions, or in regions of the parameter space where the non-convexity is known to be strong. Conversely, the hybrid solver is preferable at low resolutions or when the objective is known to be convex or sufficiently near convex. This comparison may change with an improved implementation of the hybrid solver.

In the next section, we will demonstrate the effectiveness of f_κ^∞ as a proxy for the maximum curvature in filament optimization, by solving a constrained optimization problem.

3.3. Validating QUADCOIL as a complexity metric

In this section, we present a parameter space study showing that QUADCOIL can predict the optimum normalized field error,

$$J_B \equiv \frac{1}{2} \frac{\int_{\text{plasma}} (\mathbf{B}_{\text{coil}} \cdot \mathbf{n} - B_T)^2 d^2s}{\int_{\text{plasma}} |\mathbf{B}|^2 d^2s}, \quad (25)$$

in a filament optimization for a given combination of maximum curvature κ_0 , minimum coil-coil spacing $d_{\text{cc}}^{\text{min}}$, and minimum coil-plasma spacing $d_{\text{cs}}^{\text{min}}$.

3.3.1. Filament models. Before constructing a QUADCOIL proxy, we first introduce the two filament models used in our parameter space study. In both models, all coils have equal fixed current I_{coil} based on equilibrium poloidal flux. The first model constrains the coils to a prescribed winding surface. We refer to this model as the CWS (coil on winding surfaces) model. The second model allows the filament to move freely in space, as long as the coil-plasma spacing is greater than the specified threshold $d_{\text{cs}}^{\text{min}}$. We refer to this model as the free filament model. While the free model is more realistic, the CWS model's geometry bears more similarity to the sheet current \mathbf{K} in the winding surface approach. The objective functions are, respectively:

$$J_{\text{CWS}} \equiv 1J_B + 1000J_\kappa + 1J_l + 500J_{\text{cc}}, \quad (26)$$

$$J_{\text{free}} \equiv 10J_B + 500J_\kappa + 1J_l + 100J_{\text{cc}} + 100J_{\text{cs}}. \quad (27)$$

The 5 penalty terms are:

1. J_B , the normalized squared flux, as defined in (25).
2. J_κ , the maximum curvature constraint:

$$J_\kappa = a^2 \sum_{\text{coils}} \frac{1}{2} \int_{\text{curve}} \max(\kappa - \kappa_0, 0)^2 dl. \quad (28)$$

Here, κ_0 is the target maximum curvature.

3. J_l , the maximum coil length constraint:

$$J_l = \frac{1}{a_{\text{WS}}^2} \sum_{i=1}^{N_{\text{coil}}} 0.5 \max(l_i - l_{\text{target}}, 0)^2. \quad (29)$$

Here, l_{target} is the desired threshold for the maximum coil length. a_{WS} is the effective minor radius of the winding surface.

4. J_{cc} , the coil-coil spacing constraint:

$$J \equiv \frac{1}{d_{\text{cc}}^{\text{ref}}} \sum_{i=1}^{N_{\text{coil}}} \sum_{j=1}^{i-1} \int_{\text{curve}_i} \int_{\text{curve}_j} \max(0, d_{\text{cc}}^{\text{min}} - \|\mathbf{r}_i - \mathbf{r}_j\|_2)^2 dl_j dl_i, \quad (30)$$

$$d_{\text{cc}}^{\text{ref}} \equiv \frac{2\pi(R_{\text{WS}} - a_{\text{WS}})}{N_{\text{coil}}}. \quad (31)$$

R_{WS} is the effective major radius of the winding surface. Intuitively, the reference distance $d_{\text{cc}}^{\text{ref}}$ is the effective inboard circumference of the winding surface, divided by the number of coils.

5. J_{cs} , the coil-plasma spacing constraint:

$$J_{\text{cs}} \equiv \frac{1}{a} \sum_{i=1}^{N_{\text{coil}}} \int_{\text{curve}_i} \int_{\text{surface}} \max(0, d_{\text{cs}}^{\text{min}} - \|\mathbf{r}_i - \mathbf{s}\|_2)^2 dl_i ds. \quad (32)$$

In both filament models, we choose the coil number per half-field-period closest to satisfying

$$\frac{N_{\text{coil}}}{\text{aspect ratio}} \approx 4.65. \quad (33)$$

To keep the coil-coil spacing consistent across different equilibria. The target, 4.65, is the value of this ratio measured in W7X. For faster convergence, all constraints in our filament optimization are enforced by penalty functions. As a result, they can break down when multiple conflicting constraints are present or the gradient of J_B is large.

Now, we construct the equivalent QUADCOIL problem to (26) and (27):

$$\begin{aligned} & \min_{\zeta', \theta'} f_B, \\ & \text{subject to} \\ & \max_{(\zeta', \theta')} |\mathbf{K}|^2 \leq K_{\text{max}}^2, \quad \frac{f_{\kappa}^{\infty}}{K_{\text{max}}^2} \leq \kappa_0, \quad K_{\text{max}} \equiv \frac{I_{\text{coil}}}{d_{\text{cc}}^{\text{min}}}. \end{aligned} \quad (34)$$

Here, the first constraint emulates the coil-coil spacing by limiting the maximum permitted current density. The second constraint approximates the maximum curvature limit by dividing f_{κ}^{∞} by the square of the target maximum current. This has a similar functional form to the curvature of a filament coil:

$$\kappa = \frac{\mathbf{I} \cdot \nabla \mathbf{I}}{|\mathbf{I}|^2}. \quad (35)$$

3.3.2. Numerical setup. Our parameter scan uses 38 stellarator and heliotron equilibria from the equilibrium dataset in [10, 28]. For each equilibrium, we optimize the normalized field error J_B for 125 combinations of target κ_0 , $d_{\text{cc}}^{\text{min}}$, and $d_{\text{cs}}^{\text{min}}$.

The CWS filament optimization uses the accessibility branch [29] of Simsopt [4]. We represent the (ζ', θ') coordinate of a filament constrained to a surface with a Fourier series:

$$\begin{aligned} \zeta'(t) &= \sum_{m=0}^{\text{order}} \zeta'_{c,m} \cos(mt) + \sum_{m=1}^{\text{order}} \zeta'_{s,m} \sin(mt), \\ \theta'(t) &= 2\pi t + \sum_{m=0}^{\text{order}} \theta'_{c,m} \cos(mt) + \sum_{m=1}^{\text{order}} \theta'_{s,m} \sin(mt). \end{aligned} \quad (36)$$

We first perform the optimization with $\text{order} = 2$. When it converges, we incrementally increase the order to 4, 6, and 8, and repeat the optimization using the low-order optimum as the initial condition. This is a standard practice in filament optimization. Filament optimization is non-convex. Therefore, it is very sensitive to the initial state and can be trapped at a local minimum. Directly initializing a high-order filament optimization using circular coils can lead to poor convergence. Initializing a low-order filament optimization with circular coils, and then gradually increasing the order of the curve, often improves the solve time and the optimal value of the objective function.

We represent a free filament with a Fourier series:

$$\begin{aligned} x(t) &= \sum_{m=0}^{\text{order}} x_{c,m} \cos(mt) + \sum_{m=1}^{\text{order}} x_{s,m} \sin(mt) \\ y(t) &= \sum_{m=0}^{\text{order}} y_{c,m} \cos(mt) + \sum_{m=1}^{\text{order}} y_{s,m} \sin(mt) \\ z(t) &= \sum_{m=0}^{\text{order}} z_{c,m} \cos(mt) + \sum_{m=1}^{\text{order}} z_{s,m} \sin(mt). \end{aligned} \quad (37)$$

Similarly to constrained filaments, we first perform the optimization with $\text{order} = 4$. Upon convergence, we repeat the optimization at $\text{order} = 6$, and then at 8, similar to the CWS filaments. To ensure that all filament cases have approximately equal coil spacing, we choose the coil number per half period closest to satisfying:

$$\frac{\text{total coil \#}}{\text{aspect ratio}} \approx 4.65 \text{ (the W7X value).}$$

Filament methods are sensitive to their initial states. For consistency, we initialize the constrained filaments as uniformly spaced cross sections of the winding surface $(\zeta'_{c,m}, \zeta'_{s,m}, \theta'_{c,m}, \theta'_{s,m} = 0)$, and free filaments by fitting $\text{order} = 4$ space curves to these cross sections using inverse Fourier transform.

For consistency across all equilibria, we normalize κ_0 , $d_{\text{cc}}^{\text{min}}$, $d_{\text{cs}}^{\text{min}}$, and l^{max} into 4 dimensionless parameters using their associated scale lengths:

Table 2. Values of $(k, d_{cc}^{\text{norm}}, d_{cs}^{\text{norm}}, l^{\text{norm}})$ in the grid scan, compared with measured values at NCSX and W7X. R_{ws}, a_{ws} values in this table are estimated by finding best-fit surfaces that interpolate all NCSX coils and non-planar W7X coils [4].

Scan points	NCSX values	W7X values
$k \in \{0.1, 0.2, 0.3, 0.4, 0.5\}$	0.30	0.33
$d_{cc}^{\text{norm}} \in \{0.2, 0.4, 0.6, 0.8, 1.0\}$	0.65	0.51
$d_{cs}^{\text{norm}} \in \{1.0, 1.5, 2.0, 2.5, 3.0\}$	1.73	2.45
$l^{\text{norm}} = 1.5$	1.43	1.13




Figure 6. QUADCOIL J_B v. CWS filament J_B (left) and NESCOIL J_B v. CWS filament J_B (right). The correlation is stronger in the left plot, indicating QUADCOIL is a more effective coil complexity proxy than NESCOIL.

1. $k \equiv \frac{\kappa_0^{-1}}{a}$ is the ratio between the radius of curvature and the minor radius.
2. $d_{cc}^{\text{norm}} \equiv \frac{d_{cc}^{\text{min}}}{d_{cc}^{\text{ref}}}$ is the coil-coil spacing, normalized by the reference distance d_{cc}^{ref} .
3. $d_{cs}^{\text{norm}} \equiv \frac{d_{cs}^{\text{min}}}{a}$ is the coil-plasma spacing normalized by the minor radius.
4. $l^{\text{norm}} \equiv \frac{l^{\text{max}}}{2\pi a_{ws}}$ is the total coil length normalized by an effective poloidal circumference, $2\pi a_{ws}$.

We perform grid scans in $(k, d_{cc}^{\text{norm}}, d_{cs}^{\text{norm}})$. The choice of scan points is based on the measured values from NCSX and W7X, as shown in table 2. Empirically, a coil length limit l^{norm} is often necessary for free filament optimization to converge. Since QUADCOIL cannot target this quantity, we choose the constant value of $l^{\text{norm}} = 1.5$ in all cases. Both filament models use the L-BFGS-B algorithm.

In QUADCOIL, we use 4 sin and 4 cos modes each in both the ζ' and θ' directions for $\Phi'(\zeta', \theta')$. For simplicity of implementation, in section 3.3, we will solely rely on the Shor relaxation part of QUADCOIL. As we show in the next section, QUADCOIL is often effective even when the exactness test (59) fails. As a control group, we will also compare the QUADCOIL and NESCOIL results using the same winding surfaces.

3.3.3. Results. Across all convergent combinations of equilibria, κ_0 , d_{cc}^{min} and d_{cs}^{min} , we find the J_B values measured

from both the CWS and free filament optimization correlate with the values from the QUADCOIL optima. Figure 6 shows that CWS filament coils rarely achieve a lower J_B than QUADCOIL. The correlation between J_B from the CWS filament method and NESCOIL is weaker in comparison. This suggests that QUADCOIL may serve as an empirical lower bound of the best normalized field error achievable by CWS filament coils. The spread of the points in figure 6 suggests that the correlation is stronger at higher values of J_B . This indicates that QUADCOIL is best suited for excluding equilibria that are hard to build coils for. As figure 7 shows, QUADCOIL cases have a median run-time of 22.4 s, $175 \times$ faster than the median run-time of 3865.1 s of CWS filaments.

Figure 8 shows that a similar correlation is present between the J_B values in free filament optimization and QUADCOIL. Similarly to figure 6, the correlation becomes stronger at high J_B . This correlation persists even when the solution from QUADCOIL is inexact. As figure 7 shows, the median run time of QUADCOIL is $63 \times$ faster than the median run time of 1398.7 s of the free filaments. This speed advantage allows QUADCOIL to be integrated into the equilibrium optimization loop.

In addition to validating the effectiveness of QUADCOIL as a coil complexity proxy, we also find a significant difference between the best achievable J_B , as well as their sensitivity to κ_0 , d_{cc}^{min} , and d_{cs}^{min} across all equilibria. Figure 9 compares all coil configurations in 3 example equilibria: CFQS [30], Giuliani 2022 QA [31], and LHD major radius 3.75 m [32].

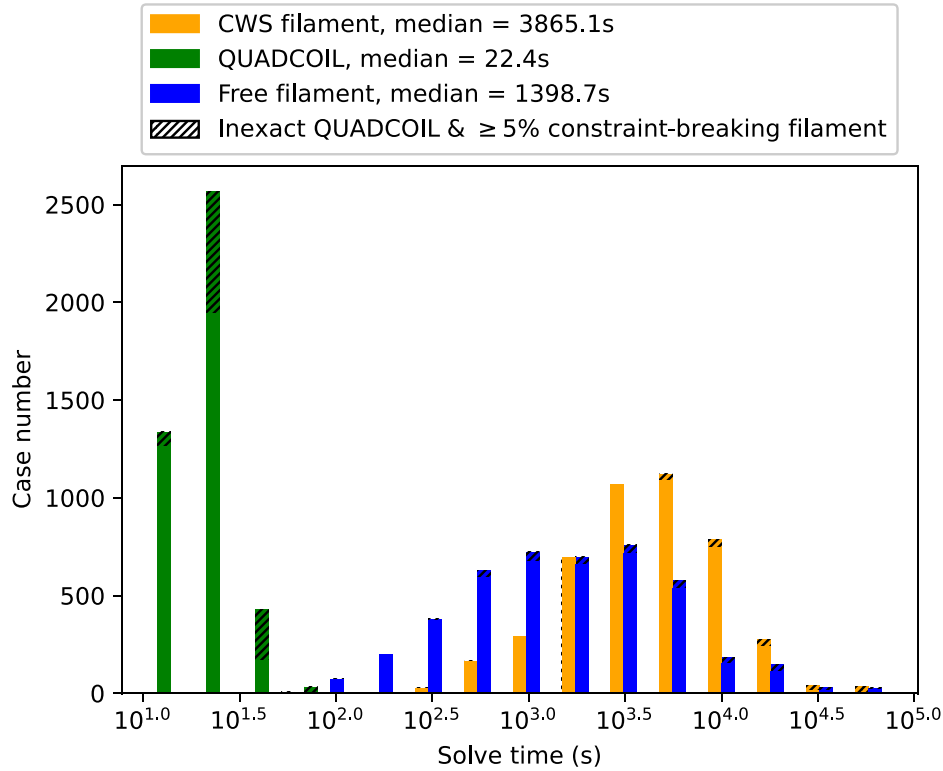


Figure 7. Run time comparison of QUADCOIL, CWS filament, and free filament. The median time of QUADCOIL is $\approx 10^2 \times$ faster than both CWS filaments and free filaments. Note that the constraints in filament optimization are enforced by penalty functions, and can be broken when multiple conflicting constraints are present, the gradient of J_B is large, or the winding surface has large principal curvatures.




Figure 8. QUADCOIL J_B v. free filament J_B (left) and NESCOIL J_B v. free filament J_B (right). The correlation is stronger in the left plot, indicating QUADCOIL is a more effective coil complexity proxy than NESCOIL.

Across all parameter combinations, the coil sets of Giuliani 2022 QA consistently achieve J_B orders of magnitude lower than coils for LHD major radius 3.75 m and CFQS. At the same time, the coil sets for Giuliani 2022 QA are also the most sensitive to changes in κ_0 , d_{cc}^{\min} , and d_{cs}^{\min} , while the coil sets for LHD are the least. Notably, across all combinations of κ_0 , d_{cc}^{\min} , and d_{cs}^{\min} , the coil sets for each equilibrium tend to form visible

‘clusters’. These observations indicate that the choice of equilibrium can have an order-of-magnitude impact on the engineering complexity of a stellarator’s coils. This strongly justifies the recent uptick in single-stage optimizations. Appendix D contains example coil sets for all three example equilibria. We will investigate the correlation between equilibrium properties and coil complexity in future studies.




Figure 9. Comparison between the QUADCOIL and free filament J_B in 3 example equilibria.

4. Conclusion

In this paper, we present QUADCOIL, a global winding surface method for stellarator optimization. We demonstrated its potential as an initial state tool and as a coil complexity proxy in the equilibrium stage. QUADCOIL hints at a deeper connection between filament and winding surface optimization and enables new parameter space studies that can help understand the relationship between equilibrium properties and coil complexity. We have also shown that different equilibrium choices can have orders of magnitude differences in coil feasibility. In future work, we will experiment with new objectives and constraints, such as local magnetic field errors, maximum dipole density, magnetic energy, and Lorentz force [19]. We will implement new representations for Φ' . A finite element basis, for example, is better suited for modeling regions with zero current. We will study the smoothness of coil complexity proxies based on QUADCOIL. By running QUADCOIL over equilibrium datasets [33–35], we may identify the equilibrium properties that lead to simple, low-force coils or dipole arrays using data-driven methods.

Acknowledgments

This work is supported by the DOE and the Simons Foundation. We are grateful for the valuable discussions with Michael C. Zarnstorff, Rory Conlin, and other members of the Simons Collaboration on Hidden Symmetries and Fusion Energy. We thank Antoine Baillo and João Pedro Biu for their kind and timely support with the Simsopt code, and John Kappel for providing the magnetic scale length dataset [28].

Appendix A. Curvature proxy $\mathbf{K}' \cdot \nabla \mathbf{K}'$

In this appendix, we discuss the construction of (12), a proxy for the maximum curvature of a coil set. $\mathbf{K}' \cdot \nabla \mathbf{K}'$ is tied to

the curvature of surface current \mathbf{K}' :

$$\kappa = \frac{\|\mathbf{K}' \cdot \nabla \mathbf{K}' \times \mathbf{K}'\|_2}{\|\mathbf{K}'\|_2^3}, \quad (38)$$

the variation of a surface current along each current path:

$$\frac{\partial_{\hat{\kappa}} \|\mathbf{K}'\|_2^2}{2} = \frac{\|\mathbf{K}' \cdot \nabla \mathbf{K}' \cdot \mathbf{K}'\|_2^2}{\|\mathbf{K}'\|_2^2}, \quad (39)$$

as well as the curvature of a filament coil with current vector \mathbf{I} :

$$\kappa = \frac{\|\mathbf{I} \cdot \nabla \mathbf{I}\|_2}{I^2}. \quad (40)$$

In practice, it is often advantageous to limit the maximum, rather than integral of an engineering metric. While $\max_{\zeta', \theta'} \|\mathbf{K}' \cdot \nabla \mathbf{K}'\|_2$ is quartic, it is bounded by f_κ^∞ via the norm equivalence principle.

$$f_\kappa^\infty \leq \max_{\zeta', \theta'} \|\mathbf{K}' \cdot \nabla \mathbf{K}'\|_2 \leq \sqrt{3} f_\kappa^\infty. \quad (41)$$

This statement is straightforward to prove. Consider a finite vector field $\mathbf{v}(\mathbf{r}) \in \mathbb{R}^3$ defined in an arbitrary, closed domain $\mathbf{r} \in \mathbb{D}$. Define:

$$\mathbf{v}_2 = \operatorname{argmax}_{\mathbb{D}} (\|\mathbf{v}\|_2), \\ \mathbf{v}_\infty = \operatorname{argmax}_{\mathbb{D}} (\|\mathbf{v}\|_\infty).$$

Then, $\forall \mathbf{v} \in \mathbb{D}$,

$$\|\mathbf{v}_2\|_2 \geq \|\mathbf{v}\|_2 \\ \|\mathbf{v}_\infty\|_\infty \geq \|\mathbf{v}\|_\infty. \quad (42)$$

Therefore,

$$\|\mathbf{v}_2\|_2 \geq \|\mathbf{v}_\infty\|_2, \\ \|\mathbf{v}_2\|_\infty \leq \|\mathbf{v}_\infty\|_\infty. \quad (43)$$

By norm equivalence,

$$\|\mathbf{v}_2\|_\infty \leq \|\mathbf{v}_2\|_2 \leq \sqrt{3} \|\mathbf{v}_2\|_\infty, \\ \|\mathbf{v}_\infty\|_\infty \leq \|\mathbf{v}_\infty\|_2 \leq \sqrt{3} \|\mathbf{v}_\infty\|_\infty. \quad (44)$$

Therefore,

$$\|\mathbf{v}_\infty\|_\infty \leq \|\mathbf{v}_2\|_2 \leq \sqrt{3} \|\mathbf{v}_\infty\|_\infty, \quad (45)$$

Although the functional form of f_κ^∞ is different from the maximum curvature of a space curve, results in section 3.3 have demonstrated its effectiveness.

Appendix B. Full expressions for \mathbf{K}' and $\mathbf{K}' \cdot \nabla \mathbf{K}'$

This appendix derives the full expression for \mathbf{K}' and $\mathbf{K}' \cdot \nabla \mathbf{K}'$ in terms of the single-valued current potential Φ' , the net poloidal current G , and the net toroidal current I , for use in the implementation.

The expression of \mathbf{K}' in terms of Φ' is:

$$\begin{aligned} \mathbf{K}' &= \frac{\mathbf{N}}{|\mathbf{N}|} \times \nabla \left[\Phi' + \frac{G\zeta'}{2\pi} + \frac{I\theta'}{2\pi} \right] \\ &= \frac{1}{|\mathbf{N}|} [(\partial_{\theta'} \mathbf{r}) \partial_{\zeta'} - (\partial_{\zeta'} \mathbf{r}) \partial_{\theta'}] \left[\Phi' + \frac{G\zeta'}{2\pi} + \frac{I\theta'}{2\pi} \right], \end{aligned} \quad (46)$$

and

$$\begin{aligned} \mathbf{K}' \cdot \nabla \mathbf{K}' &= \mathbf{K}' \cdot (\nabla \theta' \partial_{\theta'} + \nabla \zeta' \partial_{\zeta'}) \mathbf{K}' \\ &= \frac{1}{|\mathbf{N}|} [(\partial_{\theta'} \mathbf{r}) \partial_{\zeta'} - (\partial_{\zeta'} \mathbf{r}) \partial_{\theta'}] \left[\Phi' + \frac{G\zeta'}{2\pi} + \frac{I\theta'}{2\pi} \right] \\ &\quad \cdot (\nabla \theta' \partial_{\theta'} + \nabla \zeta' \partial_{\zeta'}) \mathbf{K}' = \frac{1}{|\mathbf{N}|} \left[\left(\partial_{\zeta'} \Phi' + \frac{G}{2\pi} \right) (\partial_{\theta'} \mathbf{K}') \right. \\ &\quad \left. - \left(\partial_{\theta'} \Phi' + \frac{I}{2\pi} \right) (\partial_{\zeta'} \mathbf{K}') \right]. \end{aligned} \quad (47)$$

The full expressions for $\partial_{\zeta'} \mathbf{K}'$ and $\partial_{\theta'} \mathbf{K}'$ are: (abbreviating ζ' or θ' as α)

$$\begin{aligned} \partial_{\alpha} \mathbf{K}' &= \partial_{\alpha} \left[\frac{\partial_{\theta'} \mathbf{r}}{|\mathbf{N}|} (\partial_{\zeta'} \Phi'_{\text{tot}}) - \frac{\partial_{\zeta'} \mathbf{r}}{|\mathbf{N}|} (\partial_{\theta'} \Phi'_{\text{tot}}) \right] \\ &= \left(\partial_{\alpha} \frac{\partial_{\theta'} \mathbf{r}}{|\mathbf{N}|} \right) (\partial_{\zeta'} \Phi'_{\text{tot}}) + \frac{\partial_{\theta'} \mathbf{r}}{|\mathbf{N}|} (\partial_{\alpha \zeta'} \Phi'_{\text{tot}}) \\ &\quad - \left(\partial_{\alpha} \frac{\partial_{\zeta'} \mathbf{r}}{|\mathbf{N}|} \right) (\partial_{\theta'} \Phi'_{\text{tot}}) - \frac{\partial_{\zeta'} \mathbf{r}}{|\mathbf{N}|} (\partial_{\alpha \theta'} \Phi'_{\text{tot}}) \\ &= \left(\partial_{\alpha} \frac{\partial_{\theta'} \mathbf{r}}{|\mathbf{N}|} \right) \left(\partial_{\zeta'} \Phi' + \frac{G}{2\pi} \right) + \frac{\partial_{\theta'} \mathbf{r}}{|\mathbf{N}|} (\partial_{\alpha \zeta'} \Phi') \\ &\quad - \left(\partial_{\alpha} \frac{\partial_{\zeta'} \mathbf{r}}{|\mathbf{N}|} \right) \left(\partial_{\theta'} \Phi' + \frac{I}{2\pi} \right) - \frac{\partial_{\zeta'} \mathbf{r}}{|\mathbf{N}|} (\partial_{\alpha \theta'} \Phi') \end{aligned} \quad (48)$$

The full expression for $\partial_{\alpha} \frac{\partial_{\zeta'} \mathbf{r}}{|\mathbf{N}|}$ and $\partial_{\alpha} \frac{\partial_{\theta'} \mathbf{r}}{|\mathbf{N}|}$ are: (abbreviating ζ' or θ' as β)

$$\begin{aligned} \partial_{\alpha} \frac{\partial_{\beta} \mathbf{r}}{|\mathbf{N}|} &= \frac{\partial_{\alpha \beta} \mathbf{r}}{|\mathbf{N}|} + (\partial_{\beta} \mathbf{r}) \partial_{\alpha} \left(\frac{1}{|\mathbf{N}|} \right) \\ &= \frac{\partial_{\alpha \beta} \mathbf{r}}{|\mathbf{N}|} + (\partial_{\beta} \mathbf{r}) \partial_{\alpha} (\mathbf{N} \cdot \mathbf{N})^{-1/2} \\ &= \frac{\partial_{\alpha \beta} \mathbf{r}}{|\mathbf{N}|} + (\partial_{\beta} \mathbf{r}) \partial_{\alpha} (\mathbf{N} \cdot \mathbf{N}) \left(-\frac{1}{2} \right) (\mathbf{N} \cdot \mathbf{N})^{-3/2} \\ &= \frac{\partial_{\alpha \beta} \mathbf{r}}{|\mathbf{N}|} - (\partial_{\beta} \mathbf{r}) \frac{\partial_{\alpha} (\mathbf{N} \cdot \mathbf{N})}{2|\mathbf{N}|^3}. \end{aligned}$$

Here, $\mathbf{N} \equiv \frac{\partial \mathbf{r}}{\partial \zeta'} \times \frac{\partial \mathbf{r}}{\partial \theta'}$, $\partial_{\alpha} \mathbf{N} = \frac{\partial^2 \mathbf{r}}{\partial \alpha \partial \zeta'} \times \frac{\partial \mathbf{r}}{\partial \theta'} + \frac{\partial \mathbf{r}}{\partial \zeta'} \times \frac{\partial^2 \mathbf{r}}{\partial \alpha \partial \theta'}$, $\partial_{\alpha} (\mathbf{N} \cdot \mathbf{N}) = 2\mathbf{N} \cdot \partial_{\alpha} \mathbf{N}$.

Appendix C. Relaxation procedures

In this appendix, we find a convex relaxation to (24) following a procedure common in quadratic optimization, known as Shor relaxation [20]. This procedure converts the problem into a minimization problem for a linear objective on a non-convex domain and then solves the problem in a convex superset of the domain.

Consider a scalar quadratic function of $y \equiv (\Phi', a) \in \mathbb{R}^{n_{\Phi'} + n_a}$:

$$f(\Phi', a) = y^T Q y + \alpha^T y + \beta, \quad (49)$$

$$Q \in \mathbb{R}^{(n_{\Phi'} + n_a) \times (n_{\Phi'} + n_a)}, \alpha \in \mathbb{R}^{(n_{\Phi'} + n_a)}, \beta \in \mathbb{R}.$$

To make f linear, we homogenize $f \rightarrow f_h$ by introducing an additional scalar degree of freedom, z :

$$f_h(\Phi', a, z) = y^T Q y + \alpha^T y z + \beta z^2. \quad (50)$$

The new function satisfies $f_h(\Phi', a, 1) = f(\Phi', a)$ and is bilinear in $x \equiv (\Phi', a, z) \in \mathbb{R}^{(n_{\Phi'} + n_a + 1)}$. For simplicity, define $n_x \equiv (n_{\Phi'} + n_a + 1)$. Introduce a matrix form of the unknown, $X \equiv xx^T$, then f is linear in X :

$$f_h(\Phi', a, z) = y^T Q y + \alpha^T y z + \beta c = \text{tr}(FX)$$

$$= \text{tr} \left[\begin{pmatrix} Q & \alpha \\ 0 & \beta \end{pmatrix} \begin{pmatrix} y_1 y_1 & \cdots & y_1 y_{(n_x-1)} & y_1 z \\ \vdots & \ddots & \vdots & \vdots \\ y_{(n_x-1)} y_1 & \cdots & y_{(n_x-1)} y_{(n_x-1)} & y_{n_x} z \\ y_1 z & \cdots & y_{(n_x-1)} z & z^2 \end{pmatrix} \right]$$

or, equivalently,

$$= \text{tr} \left[\begin{pmatrix} \frac{1}{2} Q + \frac{1}{2} Q^T & \frac{1}{2} \alpha \\ \frac{1}{2} \alpha^T & \beta \end{pmatrix} \begin{pmatrix} y_1 y_1 & \cdots & y_1 y_{(n_x-1)} & y_1 z \\ \vdots & \ddots & \vdots & \vdots \\ y_{(n_x-1)} y_1 & \cdots & y_{(n_x-1)} y_{(n_x-1)} & y_{n_x} z \\ y_1 z & \cdots & y_{(n_x-1)} z & z^2 \end{pmatrix} \right]$$

(some solvers work better with symmetric matrices).

$$(51)$$

Following this procedure, we can rewrite f, g_j, h_k from the QCQP in equation (9) into $\text{tr}(FX), \text{tr}(G_j X), \text{tr}(H_k X)$, and transform the equation into a linear program (LP) on a closed, non-convex, compact domain \mathbb{D} :

$$\min_{X \in \mathbb{D}, b \in \mathbb{R}^b} [\text{tr}(FX) + c^T b]. \quad (52)$$

Here, \mathbb{D} is the intersection between a polyhedron \mathbb{D}_p and a non-convex set \mathbb{D}_n :

$$\mathbb{D} = \mathbb{D}_p \cap \mathbb{D}_n \quad (53)$$

$$\mathbb{D}_p = \left\{ X \mid X \in \mathbb{R}^{n_x \times n_x}, \text{tr}(G_j X) + d_j^T b \leq e_j, \right. \\ \left. \text{tr}(H_k X) + m_k^T b = n_k, \right\} \quad (54)$$

$$\mathbb{D}_n = \left\{ xx^T \mid x \in \mathbb{R}^{n_x}, x_{n_x} = 1 \right\} \\ = \left\{ X \mid X \in \mathbb{S}_+^{n_x}, X_{n_x, n_x} = 1, \text{rank}(X) = 1 \right\}. \quad (55)$$




Figure 10. The numerical signature for exact Shor relaxation. The green dots indicate cases where the numerical exactness check, $|\lambda_2/\lambda_1| \leq \lambda_{\text{crit}}$, succeeded, and the red dot indicates cases where the check fails. From the left to the right, f_B increases, the strength of the convex penalty decreases, and eventually, Shor relaxation breaks. The transition from green to red coincides with the point where ADAM starts to outperform Shor relaxation. This validates the effectiveness of the numerical exactness check.

Here, $\mathbb{S}_+^{n_x}$ denotes the set of $n_x \times n_x$ positive semi-definite matrices. The polyhedron \mathbb{D}_p contains the information of all quadratic constraints, and \mathbb{D}_n restricts the form of the new unknown X so that the homogenized problem reduces to the original problem. Equation (52) is equivalent to an LP on the convex hull of its domain:

$$\min_{X \in \text{conv}(\mathbb{D}), b \in \mathbb{R}^{n_b}} [\text{tr}(FX) + c^T b]. \quad (56)$$

For most choices of f and G_j , $\text{conv}(\mathbb{D})$ does not have a simple expression. However, by the non-decreasing property of convex hull, $\text{conv}(\mathbb{D}) \subseteq \text{conv}(\mathbb{D}_p) \cap \text{conv}(\mathbb{D}_n)$. Because \mathbb{D} is a polyhedron, $\text{conv}(\mathbb{D}_p) = \mathbb{D}_p$. We will prove at the end of this appendix, that $\text{conv}(\mathbb{D}_n)$ is simply:

$$\text{conv}(\mathbb{D}_n) = \{X | X \in \mathbb{S}_+^{n_x}, X_{n_x, n_x} = 1\}. \quad (57)$$

The change of domain completes the convex relaxation of (9), as:

$$\begin{aligned} \min [\text{tr}(FX) + c^T b], \\ X \in \mathbb{D}_p \cap \text{conv}(\mathbb{D}_n), \\ b \in \mathbb{R}^{n_b}. \end{aligned} \quad (58)$$

It is easy to see that the solution X exactly solves (9) when $\text{rank}(X) = 1$. In practice, this gives a fast numerical test for the optimality of a Shor relaxation solution. Then, we can test its optimality (or whether $\text{rank}(X) = 1$) by comparing the ratio of the leading eigenvalue λ_1 of X with the second largest eigenvalue, λ_2 , with a small, empirical threshold λ_{crit} :

$$|\lambda_2/\lambda_1| \leq \lambda_{\text{crit}}. \quad (59)$$

In this paper, we choose $\lambda_{\text{crit}} = 10^{-3}$. Figure 10 compares the results solving (9) with Shor relaxation and the stochastic, non-convex optimization algorithm, ADAM, from all three cases in section 3.2. Note that when $|\lambda_2|$ becomes greater than 10^{-3} ,

ADAM begins to outperform Shor relaxation, indicating that the relaxation no longer yields an exact solution to (9). Now, we return to prove that the convex hull of \mathbb{D}_n ,

$$\text{conv}(\mathbb{D}_n) = \left\{ \sum_{i \in \{1, 2, \dots\}} a_i X_i \mid a_i \in [0, 1], \sum_i a_i = 1, X_i \in \mathbb{D}_n \right\}, \quad (60)$$

is:

$$\mathbb{H}_n \equiv \{X | X \in \mathbb{S}_+^{n_x}, X_{n_x, n_x} = 1\}. \quad (61)$$

To start, we first prove $\text{conv}(\mathbb{D}_n) \subseteq \mathbb{H}_n$. Substitute in the definition of $\mathbb{D}_n = \{x x^T | x \in \mathbb{R}^{n_x}, x_{n_x} = 1\}$:

$$\begin{aligned} \text{conv}(\mathbb{D}_n) &= \left\{ \sum_{i \in \{1, 2, \dots\}} a_i x_i x_i^T \mid a_i \in [0, 1], \sum_i a_i \right. \\ &= 1, x_i \in \mathbb{R}^{n_x}, (x_i)_{n_x} = 1 \left. \right\}, \end{aligned} \quad (62)$$

Then, because the last element in x_i , $(x_i)_{n_x} = 1$, the last diagonal element of $x_i x_i^T$, $(x_i x_i^T)_{n_x, n_x} = 1$. This allows us to re-write the summation, $\sum_i a_i = 1$, as:

$$\begin{aligned} \text{conv}(\mathbb{D}_n) &= \left\{ \sum_{i \in \{1, 2, \dots\}} a_i x_i x_i^T \mid a_i \in [0, 1], \sum_i a_i (x_i x_i^T)_{n_x, n_x} \right. \\ &= 1, x_i \in \mathbb{R}^{n_x}, (x_i)_{n_x} = 1 \left. \right\} \end{aligned} \quad (63)$$

Because the sum of positive semi-definite matrices is still positive semi-definite,

$$\sum_{i \in \{1, 2, \dots\}} a_i x_i x_i^T \quad (64)$$

is positive semi-definite. It is easy to see, then, that:

$$\text{conv}(\mathbb{D}_n) \subseteq \mathbb{H}_n \equiv \{X|X \in \mathbb{S}_+^{n_x}, X_{n_x, n_x} = 1\}. \quad (65)$$

This completes half of the proof.

To prove $\mathbb{H}_n \subseteq \text{conv}(\mathbb{D}_n)$, we consider the rank-1 decompositions of X . It is easy to see that X has an infinite number of orthogonal, rank-1 decompositions. Starting from the eigendecomposition of X ;

$$X = Q_+ \Lambda_+ Q_+^T. \quad (66)$$

Here, Λ_+ is a $\text{rank}(X)$, positive diagonal matrix that contains all non-zero eigenvalues of X , and Q_+ is a $\text{rank}(X) \times n_x$ matrix containing $\text{rank}(X)$ orthonormal, real eigenvectors of X . This decomposition is non-unique to an arbitrary orthogonal transformation R , allowing us to write an infinite number of rank-1 decompositions as orthogonal vectors $m_i, i = 1, \dots, \text{rank}(X)$:

$$X = \sum_{i=1}^{\text{rank}(X)} m_i m_i^T = M M^T. \quad (67)$$

Here, m_i are the rows of $M \equiv R \sqrt{\Lambda_+} Q_+ \in \mathbb{R}^{\text{rank}(X) \times n_x}$. m_i are orthogonal, but not normal.

Now, consider $(\sqrt{\Lambda_+} Q_+)_{n_x}$, the last column of $\sqrt{\Lambda_+} Q_+$. Because $X_{n_x, n_x} = 1$, $|(\sqrt{\Lambda_+} Q_+)_{n_x}|^2 = 1$. Because orthogonal transforms are length-preserving, $(m_i)_{n_x}$ must also satisfy $\sum_{i=1}^{\text{rank}(X)} (m_i)_{n_x}^2 = 1$. Choose an R that produces non-zero $(m_i)_{n_x}$ for all i . This choice of R will define a rank-1 decomposition satisfying:

$$X = \sum_{i=1}^{\text{rank}(X)} m_i m_i^T, \quad \sum_{i=1}^{\text{rank}(X)} (m_i)_{n_x}^2 = 1, \quad (68)$$

Because $\sum_{i=1}^{\text{rank}(X)} (m_i)_{n_x}^2 = 1$, and m_i are real, $(m_i)_{n_x}^2 \in (0, 1]$. Define $a_i \equiv (m_i)_{n_x}^2$, then:

$$\begin{aligned} X &= \sum_{i=1}^{\text{rank}(X)} a_i \begin{bmatrix} m_i \\ (m_i) \end{bmatrix} \begin{bmatrix} m_i \\ (m_i) \end{bmatrix}^T, \quad a_i \in (0, 1], \\ \sum_{i=1}^{\text{rank}(X)} a_i &= 1 \in \left\{ \sum_{i \in \{1, 2, \dots\}} a_i x_i x_i^T \mid a_i \in [0, 1], \sum_i^{n_{\text{obj}}} a_i \right. \\ &\quad \left. = 1, x_i \in \mathbb{R}^{n_x}, (x_i)_{n_x} = 1 \right\} \in \text{conv}(\mathbb{D}_n). \quad (69) \end{aligned}$$

This proves $\mathbb{H}_n \subseteq \text{conv}(\mathbb{D}_n)$. It may be tempting to apply (69) directly to the eigendecomposition of X . However, because there is no guarantee that the last elements of all eigenvectors are non-zero, the use of R and M is necessary.

By proving $\text{conv}(\mathbb{D}_n) \subseteq \mathbb{H}_n$ and $\mathbb{H}_n \subseteq \text{conv}(\mathbb{D}_n)$, we have proved that $\text{conv}(\mathbb{D}_n) = \mathbb{H}_n = \{X|X \in \mathbb{S}_+^{n_x}, X_{n_x, n_x} = 1\}$.

Appendix D. Random samples of free filament topology

As shown in figure 11, we choose 3 examples each from the LHD, CFQS, and Giuliani 2022 equilibria with low, intermediate, and high J_B . Figures 12, 13 and 14 show the optimum free filament coils in these examples.

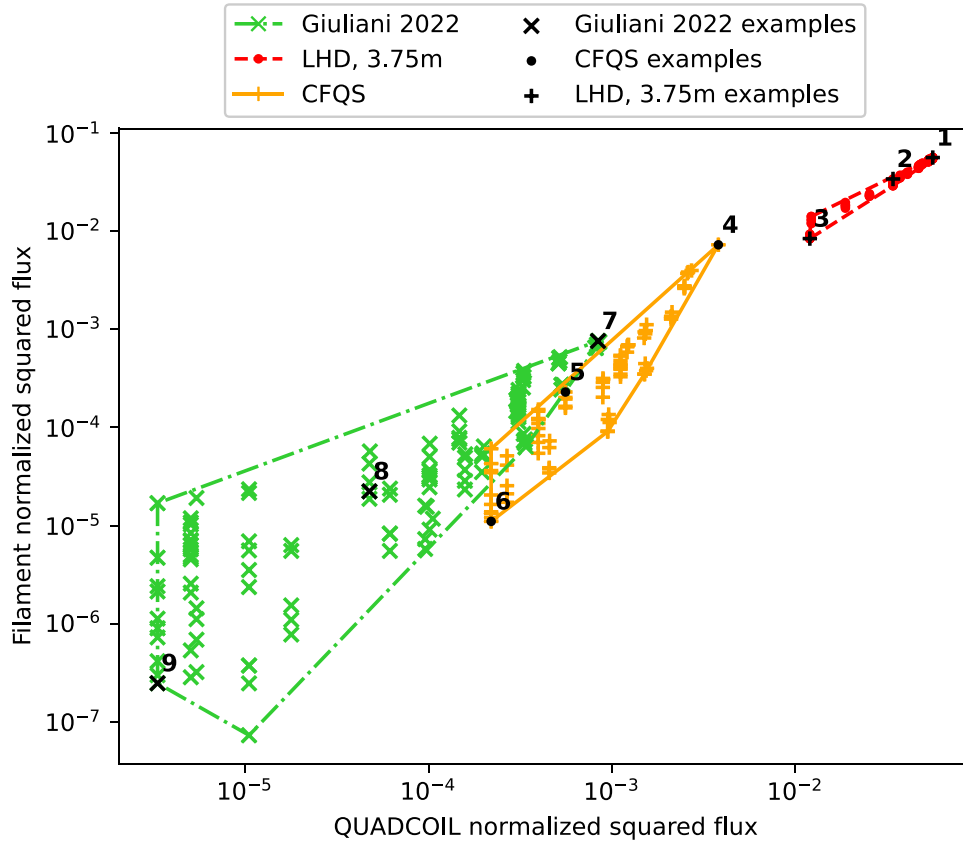


Figure 11. The locations of the nine example cases. Cases 1–3 are optimized for the LHD 3.75 m equilibrium, cases 4–6 are optimized for the CFQS equilibrium, and cases 7–9 are optimized for the Giuliani 2022 equilibrium. All cases in this figure have exact Shor relaxation and convergent free filament optimization with $\leq 5\%$ constraint breaking.




Figure 12. Example coil sets (cases 1–3 in figure 11) for the LHD 3.75 m equilibrium [32]. The normalized squared flux decreases from case 1–3.

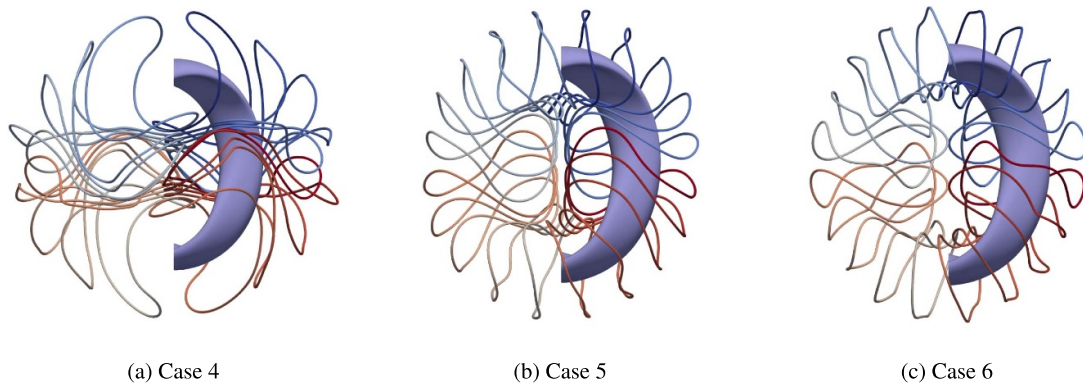


Figure 13. Example coil sets (cases 4–6 in figure 11) for the CFQS equilibrium [30]. The normalized squared flux decreases from case 4–6.




Figure 14. Example coil sets (cases 7–9 in figure 11) for the Giuliani 2022 equilibrium [31]. The normalized squared flux decreases from case 7–9.

ORCID iDs

Lanke Fu  <https://orcid.org/0000-0002-6845-387X>

Elizabeth J. Paul  <https://orcid.org/0000-0002-9355-5595>

References

- [1] Strykowski R.L., Brown T., Chrzanowski J., Cole M., Heitzenroeder P., Neilson G.H., Rej D. and Viol M. 2009 Engineering cost schedule lessons learned on NCSX 2009 23rd IEEE/NPSS Symp. on Fusion Engineering (San Diego, CA, USA, 1–5 June 2009) (IEEE) pp 1–4 (available at: <http://ieeexplore.ieee.org/document/5226449/>)
- [2] Strickler D.J., Berry L.A. and Hirshman S.P. 2002 Designing coils for compact stellarators *Fusion Sci. Technol.* **41** 107–15
- [3] Zhu C., Hudson S.R., Song Y. and Wan Y. 2018 New method to design stellarator coils without the winding surface *Nucl. Fusion* **58** 016008
- [4] Landreman M., Medasani B., Wechsung F., Giuliani A., Jorge R. and Zhu C. 2021 SIMSOPT: a flexible framework for stellarator optimization *J. Open Source Softw.* **6** 3525
- [5] Dudt D.W. and Kolemen E. 2020 DESC: a stellarator equilibrium solver *Phys. Plasmas* **27** 102513
- [6] Merkel P. 1987 Solution of stellarator boundary value problems with external currents *Nucl. Fusion* **27** 867–71
- [7] Drevlak M. 1998 *Coil Designs for a Quasi-Axially Symmetric Stellarator* (Association Euratom-CEA Cadarache) (available at: https://inis.iaea.org/search/search.aspx?orig_q=RN:31008071)
- [8] Landreman M. 2017 An improved current potential method for fast computation of stellarator coil shapes *Nucl. Fusion* **57** 046003
- [9] Drevlak M. 1998 Automated optimization of stellarator coils *Fusion Technol.* **33** 106–17
- [10] Kappel J.T., Landreman M. and Malhotra D. 2024 The magnetic gradient scale length explains why certain plasmas require close external magnetic coils *Plasma Phys. Control. Fusion* **66** 025018
- [11] Boozer A.H. 2000 Optimization of the current potential for stellarator coils *Phys. Plasmas* **7** 629–34
- [12] Pomphrey N. et al 2001 Innovations in compact stellarator coil design *Nucl. Fusion* **41** 339–47
- [13] Kaptanoglu A.A., Langlois G.P. and Landreman M. 2024 Topology optimization for inverse magnetostatics as sparse regression: application to electromagnetic coils for stellarators *Comput. Methods Appl. Mech. Eng.* **418** 116504
- [14] Giuliani A., Wechsung F., Cerfon A., Stadler G. and Landreman M. 2022 Single-stage gradient-based stellarator coil design: optimization for near-axis quasi-symmetry *J. Comput. Phys.* **459** 111147
- [15] Conlin R., Schilling J., Dudt D., Panici D., Kim P., Unalmis K.E. and Kolemen E. 2022 Free boundary stellarator equilibria and coil optimization in DESC *APS-DPP 2022* (available at: <https://meetings.aps.org/Meeting/DPP22/Session/NP11.18>)
- [16] Henneberg S.A., Hudson S.R., Pfefferlé D. and Helander P. 2021 Combined plasma–coil optimization algorithms *J. Plasma Phys.* **87** 905870226

- [17] Jorge R., Goodman A., Landreman M., Rodrigues J. and Wechsung F. 2023 Single-stage stellarator optimization: combining coils with fixed boundary equilibria *Plasma Phys. Control. Fusion* **65** 074003
- [18] Elder T.M. 2024 Three-dimensional magnetic fields: from coils to reconnection *PhD Thesis* Columbia University (available at: https://librarysearch.williams.edu/permalink/01WIL_INST/sktfin/cdi_proquest_journals_2884553613)
- [19] Robin R. and Volpe F.A. 2022 Minimization of magnetic forces on stellarator coils *Nucl. Fusion* **62** 086041
- [20] Shor N.Z. 1987 Quadratic optimization problems *Sov. J. Comput. Syst. Sci.* **25** 1–11
- [21] Agrawal A., Verschueren R., Diamond S. and Boyd S. 2018 A rewriting system for convex optimization problems *J. Control Decis.* **5** 42–60
- [22] Diamond S. and Boyd S. 2016 CVXPY: a Python-embedded modeling language for convex optimization *J. Mach. Learn. Res.* **17** 1–5
- [23] MOSEK ApS 2024 The MOSEK optimization toolbox for MATLAB manual. Version 10.1. (available at: <http://docs.mosek.com/latest/toolbox/index.html>)
- [24] Kingma D.P. and Ba J. 2017 Adam: a method for stochastic optimization (arXiv:1412.6980 [cs.LG])
- [25] Byrd R.H., Lu P., Nocedal J. and Zhu C. 1995 A limited memory algorithm for bound constrained optimization *SIAM J. Sci. Comput.* **16** 1190–208
- [26] Lanke F. 2024 Quadcoil validation dataset Zenodo <https://doi.org/10.5281/zenodo.13317583>
- [27] Nelson B.E. *et al* 2003 Design of the national compact stellarator experiment (NCSX) *Fusion Eng. Des.* **66-68** 169–74
- [28] Kappel J. 2023 Magnetic gradient scale length auxillary dataset Zenodo <https://doi.org/10.5281/zenodo.8349409>
- [29] Baillod A. 2024 simsopt (GitHub Repository) (available at: <https://github.com/hiddenSymmetries/simsopt/tree/accessibility>)
- [30] Liu H. *et al* 2018 Magnetic configuration and modular coil design for the Chinese first quasi-axisymmetric stellarator *Plasma Fusion Res.* **13** 3405067
- [31] Giuliani A., Wechsung F., Stadler G., Cerfon A. and Landreman M. 2022 Direct computation of magnetic surfaces in Boozer coordinates and coil optimization for quasisymmetry *J. Plasma Phys.* **88** 905880401
- [32] Iiyoshi A. *et al* 1999 Overview of the large helical device project *Nucl. Fusion* **39** 1245
- [33] Landreman M. 2022 Mapping the space of quasisymmetric stellarators using optimized near-axis expansion *J. Plasma Phys.* **88** 905880616
- [34] Gaur R., Dudt D., Conlin R., Panici D., Unalmis K. and Kolemen E. 2023 Exploring the database of omnigenous stellarators with DESC vol 2023 p AP01.033 (available at: <https://ui.adsabs.harvard.edu/abs/2023APS..DPPAP1033G>)
- [35] Giuliani A. 2024 QUASR: the QUAsymmetric stellarator repository Zenodo <https://doi.org/10.5281/zenodo.10581415>