Efficient computation of stellarator coils with an augmented Lagrangian optimization method

Pedro F. Gil,[1] Weiping Li,[2,3] Julianne Stratton,[3] Alan A. Kaptanoglu,[3] and Eve V. Stenson[1]
[1]*Max-Planck-Institut für Plasmaphysik, 85748 Garching, Germany*
[2]*BASIS Independent Silicon Valley, California, California 95126, USA*
[3]*Courant Institute of Mathematical Sciences, New York University, New York, New York 10012, USA*

(Received 10 October 2025; accepted 14 April 2026; published 7 August 2026)

Finding feasible coils for stellarator fusion devices is a critical challenge of realizing this concept for future power plants. Current design efforts struggle to navigate the highly nonconvex optimization landscape, spend considerable resources scanning the parameter space, and may produce suboptimal coils. In this work, we present an augmented Lagrangian approach to tackle the ill-posed problem of coil optimization. We illustrate its effectiveness and versatility by generating coils for five stellarators with very different symmetries and magnetic-field shaping. In all cases, we find Pareto-optimal coil solutions that in various ways outperform published coil sets.

DOI: 10.1103/838f-vbq3

I. INTRODUCTION

The stellarator, a toroidal magnetic confinement device, stands as a promising candidate for achieving controlled nuclear fusion. Unlike its axisymmetric counterpart, the tokamak, the stellarator relies on complex, three-dimensional magnetic fields to confine the plasma and therefore does not require the use of induced currents that can lead to plasma instabilities [1–3]. This also allows for steady-state operation, which provides a clear advantage over other technologies to become an energy source. The confining magnetic fields are generated by a system of external coils, the design of which is a critical determinant of the stellarator's performance and viability.

Early stellarators were built from relatively straightforward coils (circular and helical shapes); unfortunately, the plasma confinement turned out to be much worse than anticipated. In recent decades, there has been a resurgence in stellarator research—first in the public sector and now also in the private sector—that has been driven by advances in theory and optimization, especially with respect to designing the magnetic field itself; this has come at the cost of considerably more complicated coils. Mathematically, these coils are arbitrary three-dimensional (3D) curves in space, and current approaches to calculating them can easily produce infeasible shapes.

Traditionally, stellarator design is subdivided into two stages: stage I, where the plasma boundary shape is optimized, and stage II, where the coil shapes and currents are optimized. This procedure has been used in essentially all previous fully developed stellarator designs, including for designing currently operational experiments such as Wendelstein 7-X and HSX [4–6]. This two-stage process allows for stage-1 optimization to explore suitable physics properties for fusion reactors by solving the magnetohydrodynamic equations of force balance in a fixed boundary [7–10]. A separate stage-II optimization focuses on finding a corresponding coil set. However, the optimization of stellarator coils is a magnetostatic ill-posed inverse problem; a multitude of distinct coil configurations can, in principle, generate similar magnetic fields, yet most of these solutions are impractical from an engineering standpoint [11].

Ultimately, coil complexity and requirements for larger build tolerances resulted in delays in the construction and assembly of Wendelstein 7-X and the cancellation of NCSX as it went 70 million dollars above its baseline cost of 102 million dollars [12,13]. It has been determined that coil complexity and tolerances accounted for 66% of the cost growth of NCSX. As for Wendelstein 7-X, it took 10 years to build the experiment, suffering from several delays due to lack of engineering capacity, strict margins, and tight coil tolerances [5]. Overall, it took 10^6 hours to complete the 50 modular coils [14].

Historically, the stellarator community has explored various approaches to tackle the challenge of finding simple yet accurate coils. These have ranged from employing stochastic optimization algorithms to prevent getting stuck in local minima and to produce robust configurations, performing large-scale scans capable of exploring vast parameter spaces, investigating diverse coil topologies, such as modular, helical, or dipole coils, and even optimizing the plasma and the coils simultaneously [15–24]. Despite these varied efforts, finding robust and efficient coil optimization techniques remains a challenge.

Currently, various stellarator equilibrium concepts are also being investigated in the industry [25–29]. Most of these concepts take advantage of the various flavors of omnigeneity,

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Open access publication funded by Max Planck Society.

which is a magnetic-field property allowing for the confinement of collisionless trapped particles [30–33]. This property affects in particular alpha particles that contain a considerable portion of the energy (3.5 MeV) produced from D-T fusion reactions, essential for plasma heating. Among these flavors being pursued are the quasi-axisymmetric (QA), quasihelical symmetric (QH), and quasi-isodynamic (QI) configurations [9,34]. Even with this restriction, there is remarkable freedom in the plasma shape and plasma properties [35,36]. This represents an additional challenge for the development of a robust coil optimization scheme able to reliably deliver stellarator coil sets independently of the complexity of the magnetic configuration.

Furthermore, it is estimated that the original Stellaris configuration by the start-up Proxima Fusion will have anisotropic deformations in the order of 15 mm, arising from electromagnetic loads [25]. This can potentially result in degraded field quality, asymmetries in the heat loads of island divertors, and so forth. The immense force requirements in this example are not unique to Stellaris; rather, they are emblematic of the inevitable challenges of reactor-scale coil design for stellarators.

Finding coil configurations that simultaneously achieve high magnetic-field accuracy and adhere to these engineering requirements is a constrained and highly nonconvex optimization problem. Nonconvex constraints are very challenging, so in practice this problem is converted into a multi-objective unconstrained optimization problem with weights attached to each objective. In stellarator optimization, these weights are manually tuned to find coil solutions that seem promising, but the values of these weights are hard to choose because there is no *a priori* knowledge of their optimal values. As many as 10 to 15 weights are typically used to balance the various engineering trade-offs.

This paper describes the implementation of an augmented Lagrangian method (ALM) for coil optimization aimed at tackling this multifaceted challenge. We seek to identify best-to-date stellarator coil designs in terms of the plasma and engineering complexity metrics, while avoiding heuristic and manual tuning of the optimization weights. We demonstrate the value of this method by finding a substantial set of coil solutions that, in various ways, outperform previously published coil sets. The entirety of this work is implemented and can be reproduced in the SIMSOPT optimization suite [37]. This publication expands on a preliminary account covered in Ref. [38]. The configuration files to generate the figures and results in the present work can be found online [39].

II. OPTIMIZATION METHOD

A. Prior developments in coil optimization

Early stellarator coil optimization methods considered a single continuous toroidal surface surrounding the plasma, called a winding surface, where sheet currents could flow, and a current potential was well-defined. A solution for the current potential would be found, and then coils would be "cut" from them. This approach was used to find the Wendelstein 7-X coils and can be found in codes such as NESCOIL or REGCOIL [40,41]. More advanced winding surface methods are in development, especially for generating good initial coils and for performing single-stage optimizations [20]. However, such an approach constrains the coils to sit on a single predefined toroidal surface and therefore limits the exploration of better coil sets.

Recent efforts have introduced the concept of filamentary coils that can be described as a truncated Fourier series, allowing for the adjustment of the Fourier coefficients as the optimization is performed, and avoiding any surface restriction. This is the representation used in this paper within the SIMSOPT software framework, as described below in Cartesian coordinates for the ith coil,

$$x^i(\theta) = \sum_{m=0}^{M} x^i_{c,m}\cos(m\theta) + \sum_{m=1}^{M} x^i_{s,m}\sin(m\theta), \tag{1}$$

and similarly for y^i and z^i. Here, θ corresponds to the coil arc-length index; the order M of the decomposition is to be defined by the user. The coefficients $[x_{c,m}, x_{s,m}, y_{c,m}, y_{s,m}, z_{c,m}, z_{s,m}]$, along with the currents of the coils, constitute the $n = N[3(2M+1)+1]$ degrees of freedom for N coils of the optimization problem. Traditional coil optimization fixes or experiments with a few values of M, but stage-1 optimization often uses a "Fourier continuation" method where a series of optimization subproblems are solved with low M, and using those solutions to initialize subproblems with higher and higher M until convergence [42]. We perform a scan over M in Appendix A that indicates coil optimization routines could benefit from the same procedure. Setting $M \approx 15$ seems to yield the lowest field error, and the final coil sets for $M > 15$ tend to worsen on average. This approximate trend seems to hold for traditional coil optimization algorithms as well as the augmented Lagrangian approach proposed here.

Coil optimization can be formulated as an equality- and inequality-constrained optimization problem:

$$\begin{aligned} \underset{\boldsymbol{x}}{\text{minimize}} \quad & f(\boldsymbol{x}) \\ \text{subject to} \quad & e_j(\boldsymbol{x}) = 0, \quad j \in \mathcal{E} \\ & g_k(\boldsymbol{x}) \leqslant 0, \quad k \in \mathcal{I}. \end{aligned} \tag{2}$$

Here $f : \mathbb{R}^n \to \mathbb{R}$ is the objective function to be minimized, $e_j : \mathbb{R}^n \to \mathbb{R}$ and $g_k : \mathbb{R}^n \to \mathbb{R}$ correspond to equality and inequality constraints, $\boldsymbol{x} \in \mathbb{R}^n$ are the degrees of freedom including the Fourier coefficients representing the coil filaments and currents, $\mathcal{E}$ and $\mathcal{I}$ are the set of indices for equality constraints and inequality constraints, respectively. Consider N_c constraints in total. In previous work, the next step is to transform this very challenging constrained optimization into an unconstrained problem by constructing a weighted multi-objective function. This is the mechanism used in SIMSOPT, written in the form

$$\underset{\boldsymbol{x}}{\text{minimize}} \quad f(\boldsymbol{x}) + \sum_{l=1}^{N_c} \omega_l c_l(\boldsymbol{x}), \tag{3}$$

where c_l are the various inequality and equality constraints combined and $\omega_l \in \mathbb{R}$ for $l = 1, \ldots, N_c$ are the weights. Inequality constraints are typically recast to equality constraints by $g_l \mapsto \max(g_l - g_{\text{target}}, 0)$, since this quantity vanishes if $g_l(\boldsymbol{x})$ is below the target value. As was pointed out by Conlin

et al. for plasma optimization [43], the weights ω_l have to be set by the user, rarely represent an optimal trade-off, and therefore have to be adjusted. It is also only possible to satisfy the constraints when $\omega_l \to \infty$. Moreover, in order to circumvent the struggle of weight adjustment, large-scale scans of targets and weights are used, which can be computationally expensive and suffer restrictions to local regions of the parameter space.

B. Augmented Lagrangian formalism

To tackle these issues, the following augmented Lagrangian formalism for coil optimization is proposed with the Lagrangian,

$$\mathcal{L}_{\mathrm{A}}(\boldsymbol{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}) = f(\boldsymbol{x}) - \boldsymbol{\lambda}^{\top}\boldsymbol{c}(\boldsymbol{x}) + \tfrac{1}{2}\|\sqrt{\boldsymbol{\mu}} \circ \boldsymbol{c}(\boldsymbol{x})\|_2^2, \quad (4)$$

where

(1) $\boldsymbol{\lambda} \in \mathbb{R}^{N_c}$ is the vector of Lagrange multipliers;

(2) $\boldsymbol{\mu} \in \mathbb{R}^{N_c}$ and $\mu_l > 0$ for all l are the penalty parameters;

(3) $\boldsymbol{c} : \mathbb{R}^n \to \mathbb{R}^{N_c}$ is the vector of the original equality constraints.

The term $\boldsymbol{\lambda}^{\top}\boldsymbol{c}(\boldsymbol{x})$ is the standard Lagrange multiplier term, and $\frac{1}{2}\|\sqrt{\boldsymbol{\mu}} \circ \boldsymbol{c}(\boldsymbol{x})\|^2$ is the penalty term defining the Hadamard product between $\boldsymbol{\mu}$ and $\boldsymbol{c}(\boldsymbol{x})$. The latter penalizes constraint violations. Note that what previously were user-determined hyperparameters become now Lagrange multipliers that are automatically updated throughout the optimization, so time-consuming manual updates are avoided. Moreover, the penalty term prevents the necessity of infinite magnitudes for the ω_l, avoiding the two main disadvantages of the previous formalisms. The algorithm employed to steer and update the Lagrange multipliers is inspired by Refs. [43–45] and is here presented and adapted for clarity:

Augmented Lagrangian Algorithm.

Initialize: $p = 0, \boldsymbol{x}_0, \boldsymbol{\mu}_0 > 0, \omega_{\mathrm{tol}} > 0, \eta_{\mathrm{tol}} > 0, \tau > 1$
  $\boldsymbol{\lambda}_0 \leftarrow$ array sampled from uniform distribution in [0,1]
  $\bar{\mu}_0 \leftarrow \mathrm{Average}(\boldsymbol{\mu}_0)$
  $\omega_0 \leftarrow 1/\bar{\mu}_0$
  $\eta_0 \leftarrow 1/\bar{\mu}_0^{0.1}$
  While $\|\nabla\mathcal{L}_{\mathrm{A}}(\boldsymbol{x}_p, \boldsymbol{\lambda}_p, \boldsymbol{\mu}_p)\|_2 > \omega_{\mathrm{tol}}$ or $\|\boldsymbol{c}(\boldsymbol{x}_p)\|_\infty > \eta_{\mathrm{tol}}$
    $\boldsymbol{x}_{p+1} \leftarrow \mathrm{argmin}_{\boldsymbol{x}}\mathcal{L}_{\mathrm{A}}(\boldsymbol{x}, \boldsymbol{\lambda}_p, \boldsymbol{\mu}_p)$, solved to tol. ω_p
    if $\|\boldsymbol{c}(\boldsymbol{x}_{p+1}, \boldsymbol{\lambda}_p, \boldsymbol{\mu}_p)\|_\infty < \eta_p$ then
      $\boldsymbol{\lambda}_{p+1} \leftarrow \boldsymbol{\lambda}_p - \boldsymbol{\mu}_p \circ \boldsymbol{c}(\boldsymbol{x}_{p+1})$
      $\boldsymbol{\mu}_{p+1} \leftarrow \boldsymbol{\mu}_p$
      $\omega_{p+1} \leftarrow \max(\omega_p/\bar{\mu}_{p+1}, \omega_{\mathrm{tol}})$
      $\eta_{p+1} \leftarrow \max(\eta_p/\bar{\mu}_{p+1}, \eta_{\mathrm{tol}})$
    else:
      $\boldsymbol{\lambda}_{p+1} \leftarrow \boldsymbol{\lambda}_p$
      $\boldsymbol{\mu}_{p+1} \leftarrow \tau\boldsymbol{\mu}_p$
      $\bar{\mu}_{p+1} \leftarrow \mathrm{Average}(\boldsymbol{\mu}_p)$
      $\omega_{p+1} \leftarrow \max(1/\bar{\mu}_{p+1}, \omega_{\mathrm{tol}})$
      $\eta_{p+1} \leftarrow \max(1/\bar{\mu}_{p+1}^{0.1}, \eta_{\mathrm{tol}})$
    end if
    $p \leftarrow p + 1$
  end while

Here w_{tol} is the tolerance of the Lagrangian gradient and η_{tol} is the tolerance for the measure of constraint violation, and together they set the convergence condition. The argmin($\mathcal{L}_{\mathrm{A}}$) optimization is an internal minimization step where, for fixed Lagrange multipliers, the local optimizer L-BFGS-B [46] is used, a quasi-Newton method suitable for problems with a high number of degrees of freedom. In the context of SIMSOPT, the individual gradients of each constraint function $c_i(\boldsymbol{x})$ are calculated using a combination of analytic derivatives and automatic differentiation, where the chain rule is applied using vector-Jacobian products allowing for efficient computations.

C. Treatment of squared flux

As mentioned previously, the goal is to design feasible coils that are able to reproduce exactly the desired magnetic field for confining a fusion plasma. Therefore, the main objective of the optimization is to minimize the field error. This is achieved via the normalized squared flux, which is a measure of deviation between the coil magnetic field and the target magnetic surface,

$$f_{\mathrm{SF}} = \frac{1}{2}\frac{\int_S |\mathbf{B}\cdot\mathbf{n}|^2 \mathrm{d}S}{\int_S |\mathbf{B}|^2 \mathrm{d}S}, \quad (5)$$

where S is the plasma target surface, $\boldsymbol{n}$ is its normal vector, and $\boldsymbol{B}$ is the magnetic field generated by the coils from the Biot-Savart law,

$$\boldsymbol{B} = \frac{\mu_0}{4\pi}\sum_{i=1}^{N} I_i \int_{\Gamma_i} \frac{\mathrm{d}\boldsymbol{l}_i \times \boldsymbol{r}}{r^3}. \quad (6)$$

Here $\mathrm{d}\boldsymbol{l}_i$ is the differential arc-length element of the filament, Γ_i is the ith coil curve and $\boldsymbol{r}$ is the distance between the corresponding arc-length element and the evaluation point. When $f_{\mathrm{SF}} = 0$, the field produced by the stellarator coils is exactly tangential to the desired magnetic plasma surface. If this is achieved, then the plasma magnetohydrodynamics (MHD) equilibrium is ensured. However, computational methods cannot exactly produce $f_{\mathrm{SF}} = 0$ for the entire plasma boundary. This necessity for zero errors seems to imply that f_{SF} should serve as a candidate for the minimization function f in Eq. (4).

Instead, in this work, the function f is replaced with a dummy objective function that returns zero, and the squared flux function is moved to the constraints with all the other objectives. This is motivated by the fact that the augmented Lagrangian methods are highly adapted to minimize constraint violation. We also found that for many devices $f_{\mathrm{SF}} \lesssim 10^{-6}$ or so is sufficient for reproducing the plasma properties. In this case, we can convert the squared flux objective to $\max(f_{\mathrm{SF}} - 10^{-6}, 0)^2$. Once this constraint is satisfied, the algorithm is free to use the extra "slack" to improve engineering constraints, rather than continuing to focus on the squared flux. This is also advantageous because it does not make practical sense to exactly fit $f_{\mathrm{SF}} = 0$; there will inevitably be error fields from a host of real-world sources that will lead to deviations from the desired magnetic fields. Some precisely quasisymmetric stellarator designs, like the Landreman-Paul QA/QH [9] configurations, have fairly rapid deterioration of the degree of quasisymmetry as f_{SF} deviates from zero. Nonetheless, these configurations, which have been optimized only for aspect ratio, rotational transform, and quasisymmetry,

will not have extremely high degrees of quasisymmetry once they are turned into reactor designs. In other words, when finite-β effects, ideal MHD stability, turbulence reduction, etc. are factored into the optimization problem, the degree of quasisymmetry will decrease and subsequently we expect the sensitivity with respect to errors in f_{SF} will decrease. We believe this capability to impose f_{SF} as an inequality constraint is a substantial advantage of our approach.

D. Engineering informed coil constraints

Regarding the definitions of the usual constraints $c_i(\boldsymbol{x})$ on the optimization, we use standard metrics for coil complexity, as given in the literature [47,48]:

$$g_{\mathrm{cs}} = \sum_{i=1}^{N} \int_{\Gamma_i} \int_{S} \max(0, d_0^{cs} - \|\mathbf{r}_i - \mathbf{s}\|_2)^2 \mathrm{d}l_i \mathrm{d}S, \quad (7a)$$

$$g_{\mathrm{cc}} = \sum_{i=1}^{N} \sum_{j=1}^{i-1} \int_{\Gamma_i} \int_{\Gamma_j} \max(0, d_0^{cc} - \|\mathbf{r}_i - \mathbf{r}_j\|_2)^2 \mathrm{d}l_j \mathrm{d}l_i, \quad (7b)$$

$$g_{\mathrm{l}} = \frac{1}{2}\left[\max\left(\sum_{i=1}^{N} L_i - L_0, 0\right)\right]^2, \quad (7c)$$

$$g_{\mathrm{curv}} = \sum_{i=1}^{N} \frac{1}{2} \int_{\Gamma_i} \max(\kappa - \kappa_0, 0)^2 \mathrm{d}l, \quad (7d)$$

$$g_{\mathrm{msc}} = \frac{1}{2}\left[\max\left(\sum_{i=1}^{N} \frac{1}{L_i} \int_{\Gamma_i} \kappa^2 \mathrm{d}l - K_0, 0\right)\right]^2, \quad (7e)$$

$$c_{\mathrm{link}}(\Gamma_i, \Gamma_j) = \frac{1}{4\pi}\left|\oint_{\Gamma_i} \oint_{\Gamma_j} \frac{\mathbf{r}_i - \mathbf{r}_j}{|\mathbf{r}_i - \mathbf{r}_j|^3} (\mathrm{d}\mathbf{r}_i \times \mathrm{d}\mathbf{r}_j)\right|, \quad (7f)$$

$$g_f = \frac{1}{2} \sum_{i=1}^{N} \left(\int \max(|\boldsymbol{F}_i(l)| - F_0, 0)^2 \mathrm{d}\ell\right). \quad (7g)$$

g_{cs} is a function penalizing the minimum distance between the last closed flux surface of the plasma and the coils, preventing it from being smaller than d_0^{cs}. This can allow enough space for the positioning of a tritium-breeding blanket and shielding wall to keep the integrity of the coils and fusion device that can be damaged by high-energy neutrons. g_{cc} restrains the minimum coil-to-coil distance from being smaller than d_0^{cc}, accounting for finite dimensions of the coils as well as coil assembly and diagnostics. The length of the coils is set by a quadratic penalty g_{l} that penalizes the sum of the length of all the coils if it goes above the threshold L_0. Longer coils mean more expensive superconducting material is required to construct them. The cost function g_{curv} calculates the L_2 norm of the curvature of a coil and then penalizes the local curvature of the coil if it exceeds the threshold κ_0. Similarly, g_{msc} penalizes the mean squared curvature of a curve when it exceeds the threshold K_0, this yields smoother curves. Higher curvature coils tend to be more difficult to manufacture and assemble. To prevent configurations with interlinked coils, which would be very challenging to build in practice, the Gauss linking number c_{link} is calculated for each pair and becomes an integer $\geqslant 1$ if they are interlinked. Finally, g_f minimizes the pointwise self and mutual forces that are above some F_0 threshold representing the tolerances of the superconducting material in the coils. The details for the force estimation can be found in Hurwitz *et al.* [49]. It has become an increasing necessity to find configurations minimizing the coil-to-coil forces because the required coil currents in a reactor-scale stellarator can reach tens of Mega-Ampére values [20,48,50–52]. For the purpose of this work, the estimation of the maximum force per High-Temperature Superconducting (HTS) tape turn is done by dividing the total force value by the number of HTS turns per coil.

To determine the relative HTS material requirements for coil sets with varying number of coils per half-field period, for all the following results in this work we assume 200 HTS turns per coil for a reference coilset (i.e., 800 turns total in a four-coil solution) and recalculate the number of turns assuming fixed total current per half-field period (i.e., a three-coil solution will be assumed to have 800/3 HTS turns in each coil to keep the total current fixed). This is probably a pessimistic lower bound on the amount of HTS required for solutions with fewer coils, but it also helps to mitigate the forces in these designs. These calculations are also valid if using Low-Temperature Superconductors (LTSs).

The interest behind designing stellarators with fewer coils comes from the increased plasma access that is made available on the outboard side. Fewer coils means more space between coils for placing additional ports for plasma diagnostics, increasing the heating power, and easing remote maintenance of the device. We find in practice that often the minimum coil-coil and coil-surface distances are higher, meaning less stringent assembly tolerances and more space for a neutron-absorbing blanket. However, the price is paid in the quality of the magnetic field and a potentially higher amount of required HTS material.

Lastly, note that while we have avoided the use of manually determined weights, we do not get rid of appropriate upper and lower bounds on these objectives, as this is typically critical for avoiding undesirable coil solutions such as coils which are all topologically interlinked. Furthermore, the objectives all have natural upper and lower bounds (d_0^{cs}, d_0^{cc}, …) for feasibility that are well-motivated by physical and engineering considerations. A summary of natural upper and lower bounds for a reactor-scale stellarator, and their motivation, is provided in Table I.

E. Comparison with standard methods

In this section, we compare current filamentary methods implemented in SIMSOPT with the augmented Lagrangian approach. For this, we use Fourier continuation, where the coils are optimized with a small number of degrees of freedom, and that solution is used to initialize subsequent optimizations with more and more degrees of freedom. We start with order $M = 3$ coils and solve successive optimization problems until $M = 19$. One run therefore corresponds to 16 successive optimizations. In this example, we use the Landreman-Paul quasi-axisymmetric plasma equilibrium scaled at 1 meter major radius found in Ref. [9] and include all the constraints as specified in Eq. (7). For the standard

TABLE I. Rough estimates of upper ↑ and lower ↓ bounds for engineering metrics of a reactor-scale stellarator design.

Objective	Bound	Motivation
Normalized f_{SF}	$\lesssim 10^{-5}$	Sufficient accuracy to preserve important plasma properties.
Minimum CS distance	$\gtrsim 1.3$ m	Sufficient distance for neutron-absorbing blanket.
Minimum CC distance	$\gtrsim 0.7$ m	Sufficient distance for avoiding coil intersections and assembly difficulties from finite coil build.
Total coil length	$\lesssim$150–200 m	As small as possible to decrease the costs of purchasing superconducting material. But there must be sufficient coil length for coils to be far enough from the plasma to reduce coil-surface distance and improve squared flux. Scales with the major radius.
Max curvature κ	$\lesssim 1\ \mathrm{m}^{-1}$	Low enough local curvature to simplify geometry and manufacturing costs.
Max MSC	$\lesssim 0.5\ \mathrm{m}^{-1}$	Similar to the curvature κ, the average curvature must remain low so that the coils can be manufactured precisely without extensive cost.
Coil linking	0	Interlinking coils could greatly increase the complexity of assembly.
Max force	$\lesssim 0.5$ MN/m	Max forces must stay within the force tolerances of the HTS material, estimated for industrial REBCO-based cables [53–55].

method, as the weights are traditionally user-determined, we choose to perform 20 runs with randomly assigned weights in a small range centered around known close-to-optimal weights from previous work. We then compare these 20 runs to an identical setup but optimized using the augmented Lagrangian formalism as described in Sec. II B. The results are displayed in Fig. 1. We also repeated this comparison without continuation, i.e., direct minimization using $M = 19$, and found that almost all runs produced worse results, and the improvements using augmented Lagrangian were similar to those shown in Fig. 1.

The figure illustrates all of the resulting optimization metrics that are penalized in the constraints. These plots show that standard methods are generally unable to balance all the constraints, while the augmented Lagrangian method is able to find a solution within all the upper and lower bounds. Out of the 20 runs, runs 3, 12, and 13 violate the least number of constraints (only one) but significantly fail this constraint (a prescribed force tolerance). In contrast, the augmented Lagrangian run successfully passes all tests, without needing any weight tuning. We emphasize the strength of this result; without significant manual weight tuning, the standard methods cannot regularly produce solutions with a prescribed list of constraints like the augmented Lagrangian results.

III. RESULTS

A. Quasi-axisymmetric stellarator

The first goal for this paper is to demonstrate a novel computational framework capable of efficiently reaching and in fact improving on previously Pareto-optimal solutions. We first consider quasi-axisymmetric stellarators, which are thought to combine the benefits of both stellarators and tokamaks into one device. Like tokamaks, the coils exploit a form of toroidal invariance of the magnetic-field strength and allow for the inherent steady-state operation of stellarators. A QA optimized stellarator is, in principle, also capable of showing neoclassical transport levels comparable to axisymmetric devices, as mentioned previously. The start-up company Thea Energy is currently pursuing this concept using dipole coils for their prototype reactor EOS, which is planned to be at first a neutron source [27].

The Landreman-Paul quasi-axisymmetric plasma equilibrium scaled at 1 m major radius found in Ref. [9] was chosen in order to be able to compare with the dataset found in Ref. [52], and therefore the same optimization setup is used. This is a plasma configuration with two-field periods, aspect ratio 6 and rotational transform of 0.42, with exceptional confinement properties and a quasisymmetric field error comparable to Earth's magnetic-field amplitude.

Kaptanoglu *et al.* [48] performed 8500 stellarator coil optimizations with five coils per half-field period (hfp) in order to find trade-offs between the magnetic-field accuracy that can be achieved with the coils and the Lorentz forces that act on the coils. A similar scan was performed by Hurwitz *et al.* [52]. From the 8500 minimization runs on a supercomputer, about 41 defined a Pareto-front between the coil forces and the field accuracy, which is measured using the surface-integrated average of the normal component of the magnetic field to the target surface normalized by its amplitude: $\langle \boldsymbol{B} \cdot \boldsymbol{n} \rangle / \langle B \rangle$.

To filter suitable candidates from the 8500 configurations, multiple thresholds are applied *a posteriori* to be able to determine a valid Pareto-front and remove any undesirable coil solutions that should not be considered (i.e., interlinked coils). The thresholds include a maximum allowed curvature of 12 m^{-1}, maximum allowed coil length: 6.5 m, minimum distance between coils and plasma surface: 16.6 cm, minimum distance between coils: 8.3 cm. The number of coil optimizations that were performed in the scans is a testament to the computing power needed in previous methods to achieve optimal devices.

The initial goal of this work is to be able to map this existing Pareto front. This is done by running our augmented Lagrangian scheme and adjusting the upper linear force threshold to 9, 10, 10.5, and 12 kN/m. Each run takes around 40 minutes on a single CPU core and is stopped once the algorithm reaches the maximum number of iterations, set to 50 in this case. The quadratic flux also possesses an upper boundary, which is on purpose set to low values (10^{-15}) to force the optimizer to reach the lowest error configurations.

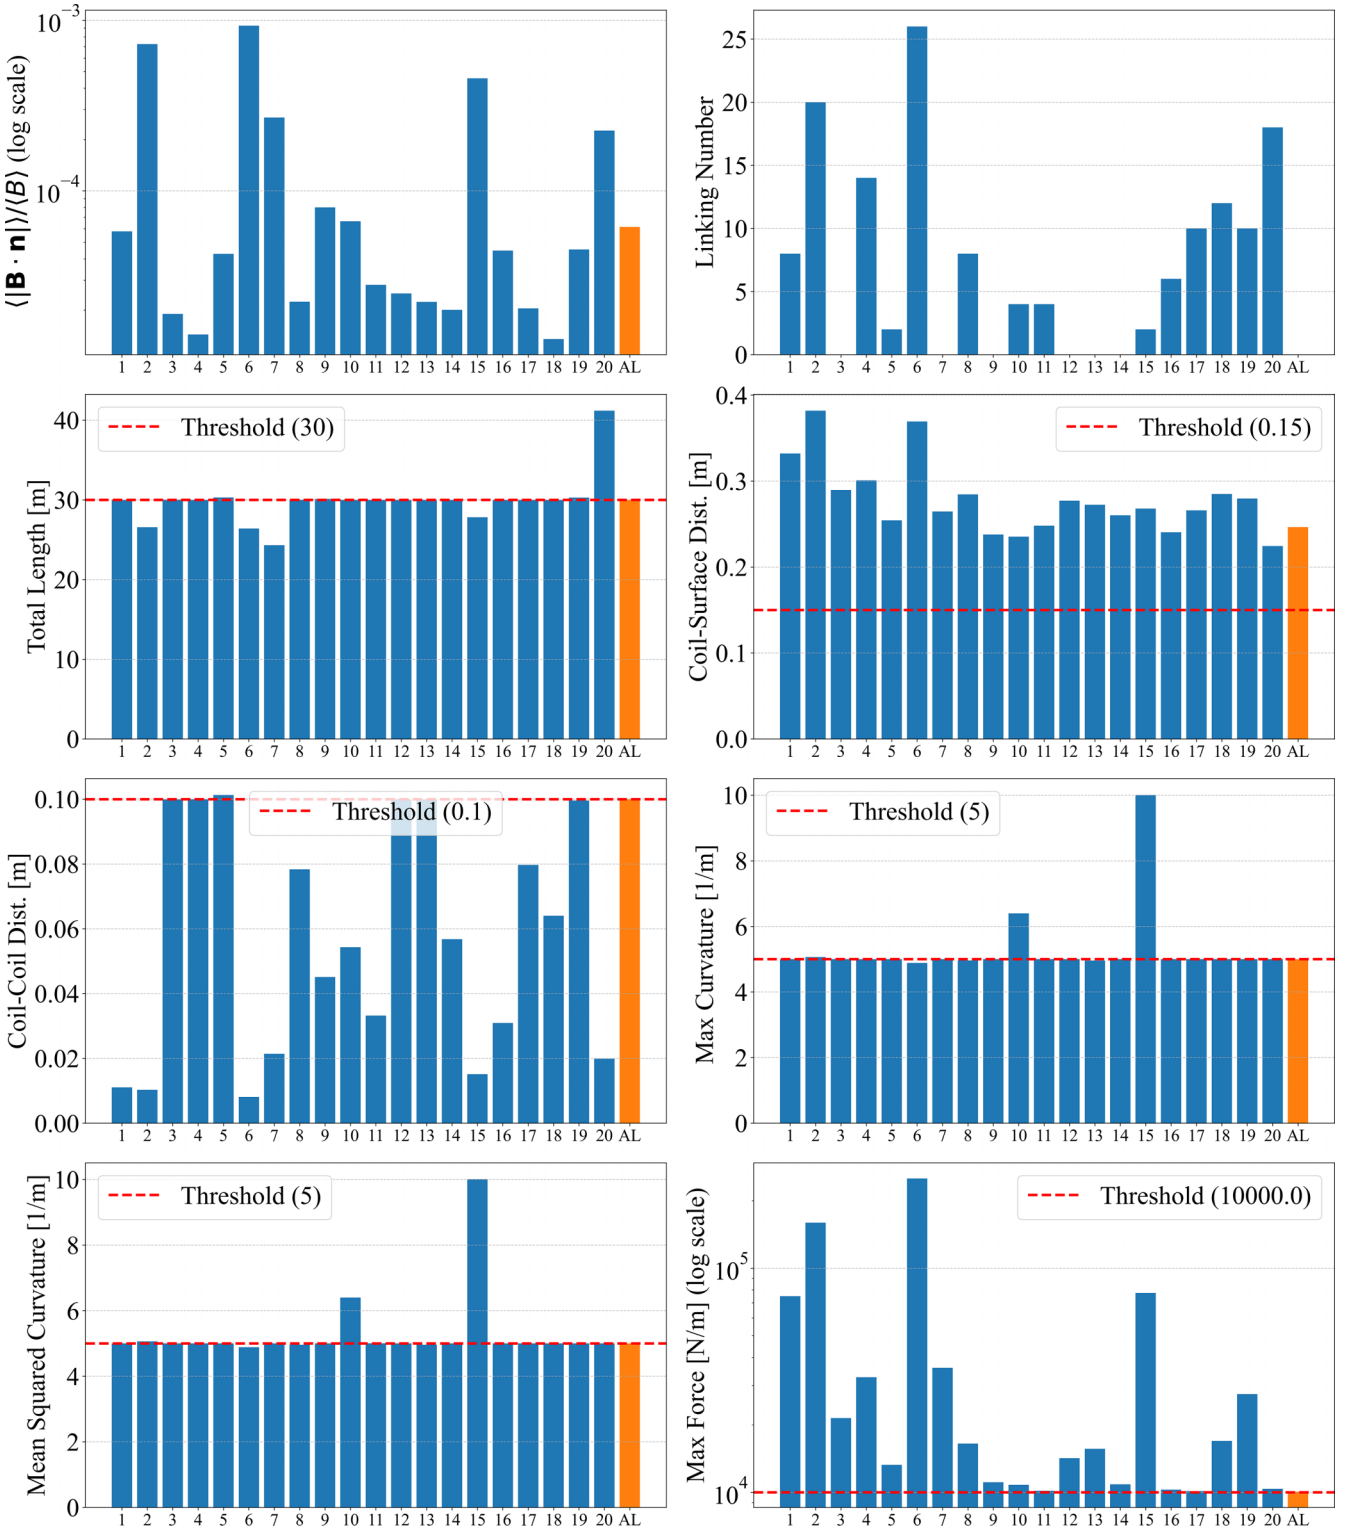


FIG. 1. Performance comparison of the augmented Lagrangian (AL) method against 20 standard coil optimization runs. The eight plots evaluate multiple targeted metrics in the optimization; the red dashed line in each plot indicates a design constraint as expressed in Eq. (7). While each of the 20 initial runs (blue bars) violates at least one constraint, the proposed augmented Lagrangian method (orange bar) is the only solution that successfully satisfies all criteria simultaneously.

All four configurations are plotted in Fig. 2, and they all sit at the Pareto-front. These results are calculated with a fixed total length of 25 m; we then try to define a new Pareto-front by relaxing this constraint and setting the total length threshold to 30 m. By running 6 different optimizations with the linear force thresholds set to 14, 12, 11, 10, 9.5, and 9 kN/m, it is possible to trace out a new limit visible in Fig. 2 with the black stars. This is a similar class of stellarators as

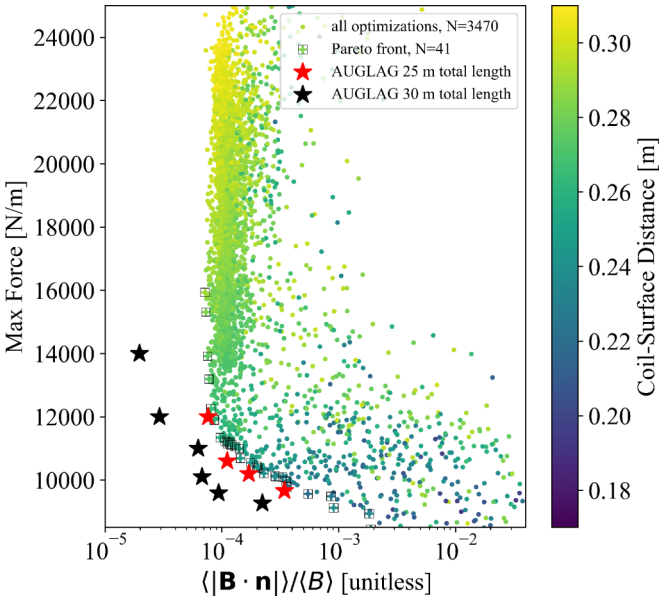


FIG. 2. Scans of QA stellarator configurations from [48] showing trade-offs between coil forces and field accuracy, with red stars placing the new configurations at the found Pareto-front and black ones outlining a new Pareto-front for configurations with larger coil length.

previously, with the main difference being that they achieve up to one order of magnitude better field accuracy, reaching a $\langle \boldsymbol{B}\cdot\boldsymbol{n}\rangle/\langle B\rangle = 1.97\times10^{-5}$ at 14 kN/m.

For the purposes of comparison with the Wechsung *et al.* [56] solutions in Table II, four-coil configurations are now scaled to reactor size with 5.7 T on-axis magnetic-field strength and 1.7 m minor radius, following ARIES-CS [57]. The total length upper bounds are set to the lengths used for the coil sets in Wechsung's solutions, and our solutions are systematically compared. For instance, the 18 m total length coils for a 1 m major radius design are indicated by L18. Scaled to a reactor, the total coil length is 182 m with a 10.13 m major radius. As far as we can tell, the Wechsung solutions are essentially Pareto-optimal. Nonetheless, without tuning the weights, we can produce an L18 solution that is essentially equivalent to the Wechsung L18 solution; the only significant difference is that the coil-coil distance and max curvature are slightly larger in our solution. Similarly, our L20 solution is essentially identical, although arguably improved by reducing the unnecessarily large coil-surface distance and increasing the minimum coil-coil distance by 0.17 m. For the L24 solution, we purposely avoid penalizing the force in order to see that only marginal improvements in the field accuracy and other metrics are available. At the same time, the max force increases by almost 50% compared with the Wechsung L24 solution. We show this result to illustrate the diminishing returns to "over-optimizing" on the field accuracy.

Lastly, at a given fixed total length, we attempted some three-coil and five-coil solutions. Five-coil solutions were found with comparable or marginally improved performance to the four-coil solutions, so we concluded that the extra coil per half-field period was not worth the inevitable reduction in plasma access. However, assuming that fairly long modular coils can be constructed in practice, an excellent solution was found with a 182 m total-length coil solution with only three coils per hfp. It replicates most of the equivalent (i.e., same total coil length) four-coil solution complexity metrics while exhibiting better field accuracy and higher coil-coil distances. It can be seen visually in Fig. 3 to provide very good plasma access. More remarkably, although it was expected that a low number of coils per hfp would reduce the degree of quasisymmetry because of increased coil ripple, the three-coil solution still manages to outperform the Wechsung L18 solution. Throughout this work, we use the two-term formulation defined in Appendix B for a quantitative measure of the quasisymmetric error displayed in Table II. Moreover, a decrease of about three orders of magnitude in the quasisymmetric error is visible as the total coil length threshold is increased. These benefits will have to be weighed against the additional coil turns in the three-coil solution that will be necessary to continue avoiding the critical current limits in the superconducting material. One plausible explanation as to why coil length impacts field accuracy is that two field-period QA configurations allow for considerable variations in

TABLE II. Comparison of magnetic-field and coil properties for our solutions and the corresponding Wechsung solutions scaled to the ARIES-CS [57] reactor with 5.7 T on axis and 1.7 m minor radius. Properties that should ideally be lowered (increased) in the design are indicated by ↓ (↑). Here, we assume a fixed total current per half-field period, taking 200 HTS turns per coil for the four-coil solution as reference, explaining the difference in the max force between the table values and Fig. 3. Top-performing metrics are highlighted in bold.

Property	Goal	Three-coil L18	Wechsung L18	L18	Wechsung L20	L20	Wechsung L24	L24
No. of coils per hfp	↓	**3**	4	4	4	4	4	4
$\langle \boldsymbol{B}\cdot\boldsymbol{n}\rangle/\langle B\rangle\times10^{-5}$	↓	85.5	91.1	93.7	32	34	**4.35**	6.0
Max $(\boldsymbol{B}\cdot\boldsymbol{n}/B)\times10^{-4}$	↓	28	34	34	13	12	**1.6**	2.9
Max Force [MN/m]	↓	**0.41**	0.44	0.43	0.47	0.46	0.55	0.85
Coil Length per hfp [m]	↓	**182**	**182**	**182**	203	203	243	240
Min CC Distance [m]	↑	**1.49**	1.34	1.44	1.14	1.31	1.0	1.0
Min CS Distance [m]	↑	2.74	2.75	2.57	2.9	2.71	3.1	**3.14**
Max κ [m^{-1}]	↓	0.53	**0.44**	0.5	0.49	0.51	0.49	0.5
Max MSC [m^{-1}]	↓	0.05	0.05	0.05	0.05	0.05	0.05	0.05
Quasisym. Error $\times10^{-6}$	↓	2000	2100	1900	190	260	2.4	10.7
HTS length [km]	↓	194	**146**	**146**	162	162	194	192

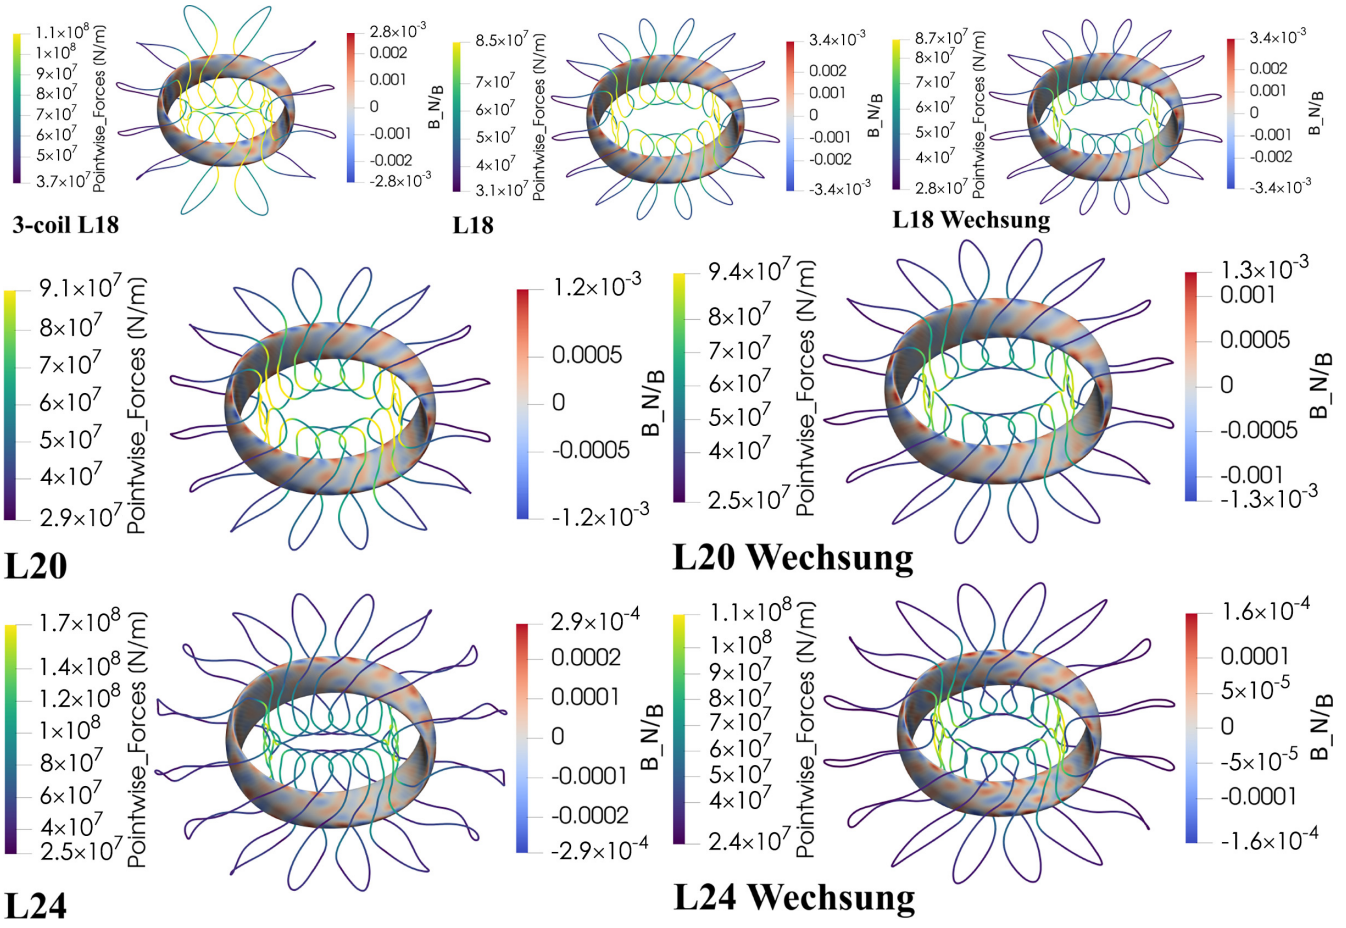


FIG. 3. (Top, left to right) Three-coil solution, four-coil solution, and four-coil Wechsung solution, all with the same total coil length. (Middle, left to right) Four-coil solution with augmented Lagrangian versus the Wechsung solution obtained at the same total coil length. (Bottom, left to right) Same as the middle panel, but with increased total coil length. Forces appear very large because they are reported here before dividing by the number of turns.

coil length. This is especially true on the inboard side, while other configurations do not (i.e., a four-field-period QH has relatively low space inside the corners), and increasing the length can effectively help to compensate coil ripple. This is, however, only valid at comparable plasma aspect ratios.

For completeness, the Boozer plot, which displays qualitatively the departure from quasisymmetry together with the Poincaré plots showing the nested flux surface structure of the magnetic equilibrium, is shown in Appendix C. Straight lines indicate that very low levels of quasisymmetric error are achieved. We also show in Appendix D that the three-coil solution found here keeps fast-particle losses to roughly 6%; this is excellent performance, especially given that this solution primarily addresses improving engineering criteria such as plasma access, forces, total length, and coil-coil distances.

Furthermore, the reason why, out of the 8500 runs performed in previous work, only 41 defined a Pareto-front, lies in the fact that the adjustment of the weights, but also the arbitrary truncation of thresholds, modifies the stellarator optimization landscape in such a way that it prevents convergence toward good configurations. In contrast, our proposed method takes care of the weight adjustment dynamically throughout the optimization, and informs their update based on the penalty violation of each of the constraints.

B. Quasihelical symmetric stellarator

In a second optimization demonstration, we target the plasma boundary of the Landreman-Paul quasihelically symmetric (QH) configuration [9]. Like quasi-axisymmetric plasmas, QH equilibria are designed to confine collisionless trapped particles while also having reduced alpha-particle losses. The QH configuration has only been effectively built once with the Helically Symmetric eXperiment (HSX) [6], which is explored in Sec. III E. One of the reasons QH symmetry and more generally quasisymmetric stellarators are of interest for a potential fusion reactor lies in the prediction of shear flows in the plasma, allowing for the suppression of microturbulence [58]. However, a four-field-period QH configuration is considered to be more demanding from a coil optimization standpoint since the plasma contains more shaping than the QA configuration shown previously. The objective is to reproduce particle confinement properties of this configuration to high enough accuracy. Therefore, the coils obtained here are compared and optimized using the same setup as in Ref. [59], where the main purpose was to obtain a suitable coil configuration at reactor scale. The major radius of these configurations is approximately 13 m, the aspect ratio is approximately 8, and the average rotational transform is 1.24.

TABLE III. Comparison of coil properties for the three different QH configurations, assuming 200 turns per coil for the max force calculations. Properties that should ideally be lowered (increased) in the design are indicated by ↓ (↑). Top-performing metrics are highlighted in bold.

Property	Goal	No. 1	No. 2	No. 3	Wiedman
No. of coils per hfp	↓	**4**	5	5	5
$\langle \boldsymbol{B}\cdot\boldsymbol{n}\rangle/\langle B\rangle \times 10^{-4}$	↓	6.2	6.2	122	**6.1**
Max $(\boldsymbol{B}\cdot\boldsymbol{n}/B)\times 10^{-3}$	↓	2.7	**2.6**	41.2	3.1
Max Force [MN/m]	↓	**0.2**	0.36	0.33	0.31
Coil Length per hfp [m]	↓	**160**	177.8	**160**	177.8
Min CS Distance [m]	↑	1.91	1.59	**2.0**	1.62
Min CC Distance [m]	↑	0.8	1.09	**1.20**	1.09
Max κ [m^{-1}]	↓	1	0.77	**0.39**	0.81
Max MSC [m^{-1}]	↓	0.1	0.08	**0.05**	0.079
Quasisym. Error $\times 10^{-3}$	↓	5.2	2.3	1602	**2.2**
HTS length [km]	↓	320	284.5	**256.0**	284.5

In the original work in Wiedman *et al.* [59], five coils per hfp were chosen as it was considered that fewer coils did not yield high enough accuracy to reproduce the plasma equilibrium at the coil-to-coil and coil-to-surface distances required. Here, we show that not only is it possible to reproduce a similar configuration as in Ref. [59], but it is also possible to both improve the engineering metrics and reduce the number of coils to four per half-field period. The total coil length is also reduced by slightly relaxing other physical requirements. The results for all three optimizations are displayed in Table III. Both the average field accuracy and maximum field error are comparable to the original work for both the No. 1 and No. 2 configurations. The stellarators are shown in Fig. 4. Except for 0.2 m reduced minimum coil-coil distance and more HTS material, the four-coil solution is as good, or better, in every engineering metric as the Wiedman solution.

To further verify that these coils generate magnetic fields that fulfill physics requirements, we plot the Boozer plot of the outermost closed magnetic surface together with the Poincaré cross sections of the equilibrium, shown in Appendix C. As expected, the magnetic-field amplitude contours appear straight, with the low-field regions displaying higher levels of coil ripple than as reported by Wiedman *et al.*; however, this is consistent with the fact that the configuration has one less coil per half-field period than previously. It is also possible to reveal the nested structure of the QH equilibrium without any apparent magnetic island chains, capturing the high degree of shaping of the equilibrium.

Confinement of collisionless trapped particles is also strongly suggested by the low quasisymmetric error present in configurations No. 1 and No. 2. Note that this is achieved without necessarily breaching the thresholds and activating the constraints in the optimization. In particular, this is visible in Table III for the coil-surface distance (lower bound set to 1.5 m) and the maximum curvature (upper bound set to 1 m^{-1}) for the No. 1 and No. 2 configurations, respectively. This shows that it is possible to set loose upper and lower bounds and still obtain feasible solutions that do not necessarily achieve the bounds.

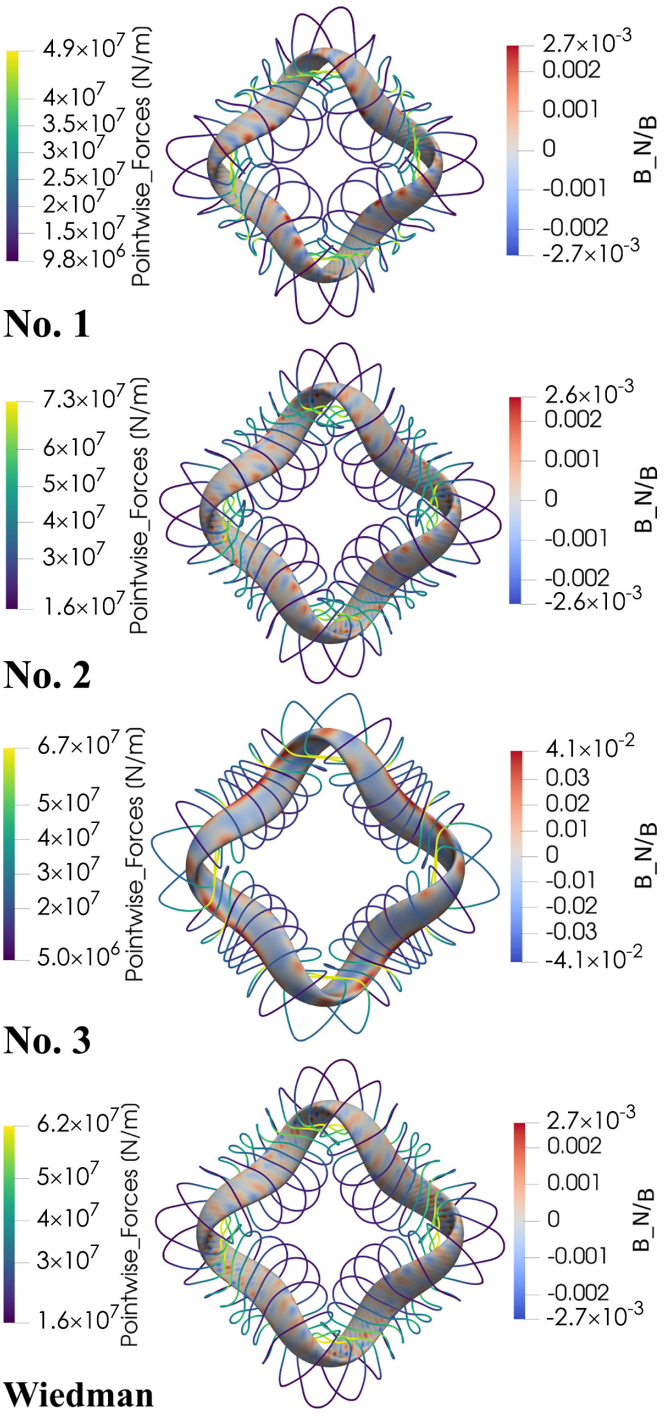


FIG. 4. QH configurations offering new trade-offs between physics and engineering. (No. 1) The stellarator coil configuration for the four-coil case. (No. 2) The stellarator coil configuration for the five-coil case. (No. 3) The stellarator coil configuration designed for optimal engineering properties. (Wiedman) Original five coils per half-field period QH coilset optimized by Wiedman *et al.* [59]. Forces appear very large because they are reported here before dividing by the number of turns.

Finally, the goal of configuration No. 3 is to show that by relaxing the field accuracy target our approach is able to improve every engineering metric that is optimized. As shown

in Table III, while the field accuracy is in the order of 1%, configuration No. 3 achieves all the best engineering properties, except for the forces, which in this case were not directly targeted during optimization. Notably, there is a 38 cm increase of coil-to-surface distance, a 17.8 m decrease in the coil length per hfp, and a 11 cm increase in the coil-to-coil distance. This translates into a 10% cost reduction in HTS tape, while having considerably simpler coils with 52% less local curvature and 37% less mean squared curvature. The simplicity of the coils is visible in Fig. 4. The maximum force also remains comparable to the Wiedman case. However, the quasisymmetric error reaches values that are three orders of magnitude higher than the other configurations. Although the field accuracy is within reasonable bounds ($\langle \boldsymbol{B}\cdot\boldsymbol{n}\rangle/\langle B\rangle \approx 1\%$), it is larger by a factor of 20 than the Wiedman configuration, and the rapid deterioration of quasisymmetry comes from high sensitivity of the equilibrium to field errors from the coils. As was mentioned previously, the Landreman-Paul configuration was targeted to reach minimal levels of quasisymmetric errors. Complementary, single-stage methods can be employed to improve the compatibility between equilibrium and coils. In Sec. III E, we present results in the same spirit as for this configuration but for the fully manufactured and assembled QH HSX stellarator.

C. Stellaris: Quasi-isodynamic stellarator

A third demonstration of the improvement brought by our method is shown on a Quasi-Isodynamic (QI) stellarator configuration [34] that was recently published by the fusion start-up Proxima Fusion [25] to serve as prototype for a fusion power plant. This builds upon the promising experimental results of the Wendelstein 7-X experiment, an optimized near-QI device [60]. The interest behind this class of stellarators lies in their physics properties and recent optimization reported in Ref. [34], where they appear MHD stable, as a consequence of the maximum-$\mathcal{J}$ property described in Refs. [61,62]. This property is also predicted to mitigate turbulence from the trapped electron mode (TEM) [3,63,64]. The same class of stellarators is being pursued by other companies such as Type One Energy [26]. Moreover, stellarators such as Stellaris display low small collisional thermal transport, bootstrap current, and fast-ion losses, which can have a great impact on plasma heating, stability, and divertor operation [3,62,65]. Finally, there is also evidence suggesting that ion-temperature gradient (ITG) turbulent heat fluxes of stellarators of this class are below that of the W7-X stellarator [34]. However, for past QI configurations it has been challenging to find buildable and accurate coils such as W7-X. On top of that, fine-tailored field line curvature imposing the maximum-$\mathcal{J}$ property or MHD stability renders the task even more complex. Proxima Fusion published a four-field-period QI configuration of aspect ratio 9.8 with major radius 12.7 m, minor radius 1.3 m, a plasma volume of 425 m^3 (about half the volume of ITER), an axis-averaged magnetic strength of 9 T, and is supposed to deliver 2700 MW of peak fusion power. The coilset found by Proxima consists of six coils per half-field period. While this allows for more field accuracy, it also comes at the cost that by accounting for the finite dimensions of the winding pack, the minimum casing-to-casing distance reaches 11.4 mm, and assuming 200 turns per coil, peak HTSs produce a linear force on the filamentary coils of 0.9 MN/m. This value is at the peak material force tolerances and limits the possibility for structural supports. These very tight engineering constraints are typical for reactor-scale stellarators.

Here three alternative configurations are presented, one with five coils, two with six coils per half-field period. In the former case, setting similar thresholds as the characteristics of Proxima's coilset except for the intercoil distance, yields a configuration that improves on the magnetic-field error average by 25%, as well as giving an additional 13 cm of space between coils, which allows for better handling of eventual inaccuracies from manufacturing, assembly, thermal contraction, or magnetic forces, which appear 5% larger. The coilset is displayed in Fig. 5. The main parameters are summarized in Table IV. Taking 200 turns per coil in the six-coil solutions as a reference, and fixed total current per half-field period, it is possible to extrapolate the amount of HTS length required for the two solutions. While the two solutions come at a slightly higher cost of HTS tape, these configurations can enable extra access to the plasma and facilitate reactor maintenance, which is a cost driver. Moreover, the six-coil case No. 2 improves on both the five-coil configuration and Proxima's coils with a field error that is cut by around half, linear forces on the HTS turns that are 29% lower, and similar distances between coils and between the plasma surface and the conductor as Proxima. Note that this is, however, a conservative estimate of the amount of HTS required because it assumes that the tape is always operated at the limit of arbitrarily set engineering requirements for the critical current. To achieve the reduced force configuration No. 2, the curvature and length constraints were relaxed, allowing the total length of the conductors to reach 145 m. The corresponding Boozer plots, which serve to verify the departure from the target isodynamicity, are shown in Appendix C.

Furthermore, configuration No. 3 takes advantage of the new threshold formulation for the quadratic flux by relaxing it while making every other constraint stricter. This yields a configuration with field errors four times higher than Stellaris but improves essentially every engineering metric. This is of special interest since it is unclear if achieving accuracies of $\langle \boldsymbol{B}\cdot\boldsymbol{n}\rangle/\langle B\rangle \approx 10^{-3}$ is at all possible at a reactor scale without auxiliary coils. Therefore, we choose to optimize for a coilset with higher errors but with, for example, more than 22 cm of spacing between coils, leaving space for error-correction coils. This is achieved while having simpler geometries with a 39% decrease in local curvature, a 35% decrease in mean-squared curvature, and 8 m/turn shorter coils, giving a total 6% decrease in HTS cost. Similarly to configuration No. 2, forces are also reduced by 29%. Most importantly, despite the substantial engineering improvements of configuration No. 3, fast-particle tracing comparisons in Appendix D show that total losses remain below 4%. As shown in Table IV, the Stellaris configuration is improved in essentially all aspects by the three different coil solutions found here. Notably, configuration No. 3 is able to improve all the engineering metrics, including 29% less force and a 6% reduction in estimated HTS cost. Nonetheless, final decisions about coil

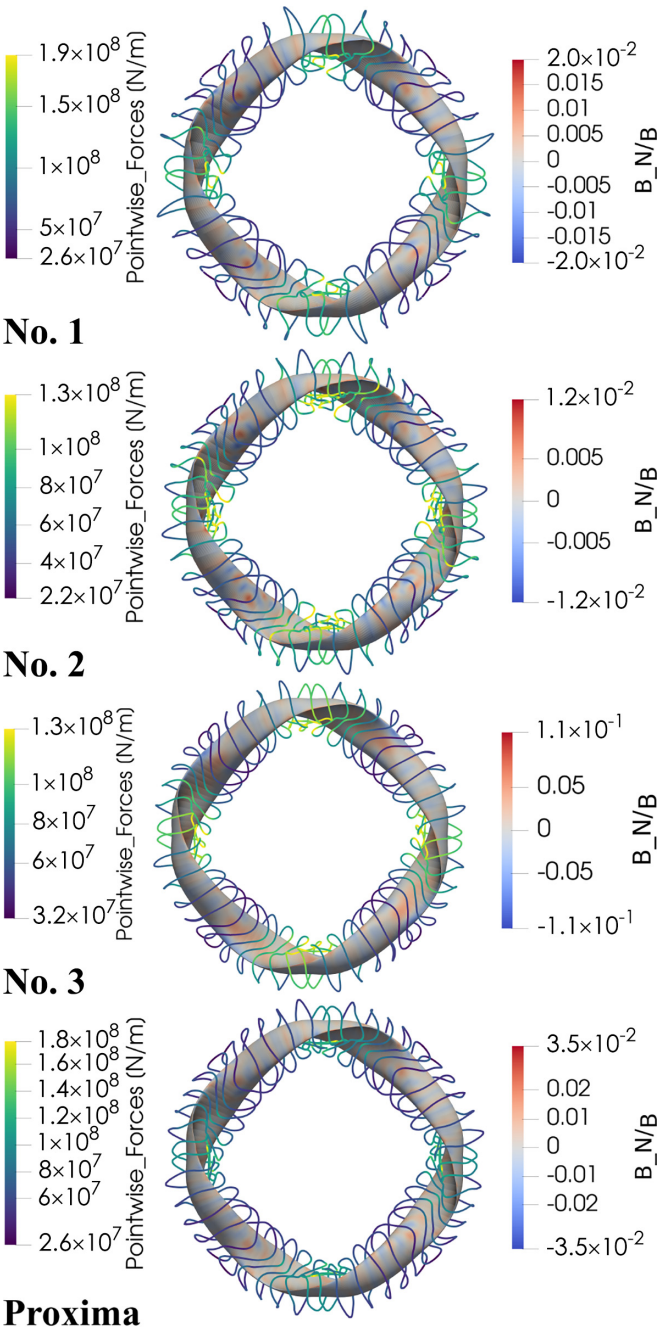


FIG. 5. New alternative coil configurations found for the Stellaris fusion power plant concept. (No. 1) First possible configuration for the Stellaris configuration with five coils per half-field period, same field errors. (No. 2) Second possible configuration with similar parameters but forces are reduced by 29%. (No. 3) Third possible configuration designed for optimal engineering properties. (Proxima) Proxima's configuration with six coils per hfp, recently published in Ref. [25]. Forces appear very large because they are reported here before dividing by the number of turns.

and HTS stack design require full finite-element modeling of the coils and structural supports beyond the scope of this work.

TABLE IV. Comparison of magnetic field and coil properties for the three different Stellaris configurations. No. 1, No. 2, and No. 3 are results from this work; the Proxima configuration was optimized by the start-up Proxima Fusion. Forces were scaled by assuming 200 turns per coil in the six-coil solution and a fixed total coil current per half-field period. Top-performing metrics are highlighted in bold.

Property	Goal	No. 1	No. 2	No. 3	Proxima
No. of coils per hfp	↓	**5**	6	6	6
$\langle \boldsymbol{B}\cdot\boldsymbol{n}\rangle/\langle B\rangle \times 10^{-3}$	↓	3.0	**2.3**	14.7	4.0
Max $(\boldsymbol{B}\cdot\boldsymbol{n}/B)\times 10^{-2}$	↓	2.2	**1.6**	11.4	3.5
Max Force [MN/m]	↓	0.79	**0.64**	**0.64**	0.9
Coil Length per hfp [m]	↓	**130**	145	**130**	138
Min CC Distance [m]	↑	0.8	0.7	**0.89**	0.67
Min CS Distance [m]	↑	1.38	1.3	**1.40**	1.37
Max κ [m^{-1}]	↓	1.57	1.6	**0.96**	1.57
Max MSC [m^{-1}]	↓	0.3	0.35	**0.22**	0.34
HTS Length [km]	↓	250	232	**208**	221

D. Wendelstein 7-X (W7-X)

So far, the configurations described correspond to devices that have been conceptualized but not yet built. In this and the next sections, we show that it is possible to improve the coil designs of stellarators that have been fully manufactured and assembled and have already delivered experimental results. The predecessor of Wendelstein 7-X, Wendelstein 7-AS, was the first stellarator to successfully demonstrate that the two-stage approach to stellarator design, as previously described, is able to generate a feasible reactor [11]. To bring forward the stellarator program, the main goals of W7-X were to demonstrate the generation of nested magnetic surfaces with modular superconducting coils, enhanced fast-particle confinement, reduced parallel currents, minimal neoclassical transport at low collisionality, and MHD stability up to an average plasma beta of 5% [66]. Wendelstein 7-X is currently the largest operating optimized stellarator in the world, with a major radius of 5.5 m and an aspect ratio of 10, being at the forefront of stellarator research and fusion science generally. The first experimental campaign was performed in 2015, using NbTi superconducting coils to generate a quasi-isodynamic magnetic field with 2.5 T on axis. In addition to its nonplanar coils, W7-X has a set of planar coils necessary for flexibility in the magnetic configuration. As mentioned in the introduction, the W7-X manufacturing and assembly process presented various delays due to tight coil tolerances. In practice, this meant that, for example, "six magnets were returned to the manufacturers three times for repair, and six others were returned twice for repair before passing the tests successfully" [67]. This resulted in additional costs.

Three different configurations are presented as options for W7-X in Fig. 6. As illustrated in Table V, all the configurations are optimized to possess lower curvatures than the original W7-X coilset. Here, the last closed flux surface of W7-X that does not contain coil ripple is used as the target. The results are visible in the coils shown in Fig. 6, as all the newly optimized coils appear smoother and with less local curvature than the W7-X ones. Note that all configurations possess comparable lengths to the W7-X coils. This freedom

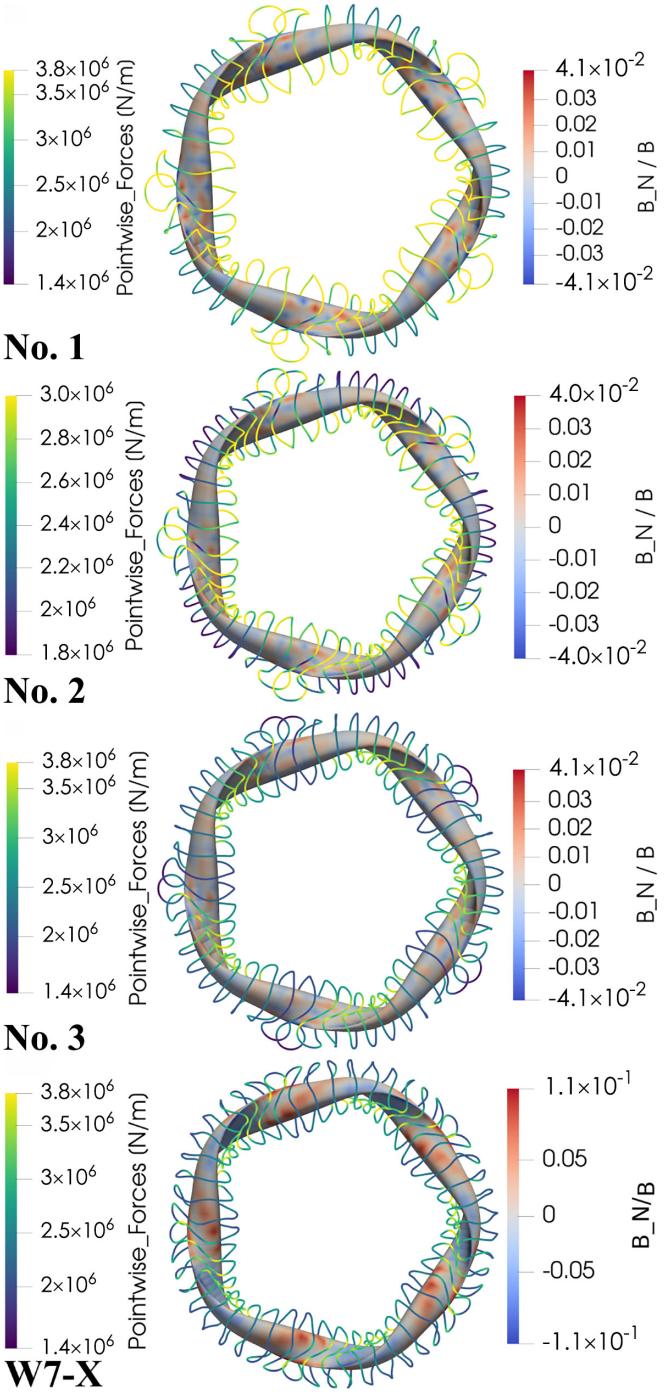


FIG. 6. Newly optimized W7-X coil sets with the augmented Lagrangian formalism. (No. 1) First alternative configuration for the W7-X stellarator with four coils per half-field period, improved field errors, simpler coils. (No. 2) Second alternative coilset for the W7-X stellarator with improved field accuracy, 20% less force, longer coils, but lower curvature. (No. 3) Third alternative coilset for W7-X with improved field error, similar coil length, same force, and simpler coils. (No. 4) Original W7-X coilset extracted from the CAD design with five coils per half-field period. Forces appear very large because they are reported here before dividing by the number of turns.

is purposely imposed for this study, as the goal is to explore options to the current W7-X reactor coils and not to replicate them.

TABLE V. Comparison of magnetic-field and coil properties for the four different W7-X configurations. No. 1, No. 2, and No. 3 are results from this work; the W7-X configuration was taken from the original CAD files and adapted as single filaments for the W7-X standard configuration (without additional planar coils). Here, the max force and HTS length are computed with 200 HTS turns per coil using the five-coil solution and fixed total current per half-field period. Top-performing metrics are highlighted in bold.

Property	Goal	No. 1	No. 2	No. 3	W7-X
No. of coils per hfp	↓	**4**	5	5	5
$\langle \boldsymbol{B}\cdot\boldsymbol{n}\rangle/\langle B\rangle \times 10^{-3}$	↓	10	6.7	**6.6**	29
Max $(\boldsymbol{B}\cdot\boldsymbol{n}/B)\times 10^{-2}$	↓	5.5	**4.1**	4.2	11
Max Force [kN/m]	↓	**15**	**15**	19	19
Coil Length per hfp [m]	↓	**38**	45	43	42.6
Min CC Distance [m]	↑	**0.3**	0.27	0.26	0.28
Min CS Distance [m]	↑	0.32	0.3	0.34	**0.37**
Max κ [m^{-1}]	↓	2.51	**2.0**	**2.0**	2.56
Max MSC [m^{-1}]	↓	**1.5**	**1.5**	**1.5**	1.96
HTS Length [km]	↓	95	90	86	**85.2**

Configuration No. 1 presents substantial advantages over the original W7-X coilset; the full reactor requires 40 coils instead of 50 but still provides better field accuracy, reduced total coil length, increased minimum coil-coil distance, and reduced curvatures. In every metric, it is improved over the original design, except for a 5 cm reduced minimum plasma-coil distance. Access to the plasma is substantially increased, as can be seen visually in Fig. 6. The reduced number of coils can help to drive costs down by reducing material, reducing assembly difficulties, and increasing the amount of plasma access for diagnostics, which can become a bottleneck in the operation of W7-X. Configuration No. 2 is targeted to reduce the electromagnetic forces compared with the original design. By allowing only an increase of 5% of the total coil length, slightly closer coils to each other and to the plasma, it is possible to achieve a reduction of about 20% of the force load on the support structure and coils of W7-X. Note that it has been reported by Pedersen *et al.* [68] that the rotational transform profile and island chain structure are modified due to deformations in the coil shapes resulting from increasing electromagnetic forces. Additionally, this new configuration comes with an improvement of about an order of magnitude in field accuracy. Configuration No. 3 is also an exciting find. With very tiny changes of 0.4 m more total coil length, 0.02 m reduction to minimum coil-coil distance, and 0.03 m reduction to minimum coil-surface distance, the squared flux and curvatures are dramatically improved, without change to the max forces.

The degree of omnigeneity for all coil sets displayed by the magnetic-field strength contours is plotted in Appendix C. All W7-X alternatives show contours following the target field. As shown in Fig. 6, the No. 2 and No. 3 alternatives possess four visibly simple coils out of five, with only one coil type standing out as longer. It also appears that lower overall forces can be achieved by relaxing the length constraint on an individual coil type. To our current knowledge and given the current parameters of the W7-X coils displayed in Table V, any of these three alternatives of W7-X could constitute viable experimental devices.

TABLE VI. Comparison of magnetic-field and coil properties for the four different HSX configurations. No. 1, No. 2, and No. 3 are results from this work; the HSX configuration was taken from the original CAD files and adapted as single filaments for the HSX standard Quasi-Helical Symmetry QHS configuration. Here, the max force is the force on 200 turns of HTS per coil in the reference six-coil solution. Top-performing metrics are highlighted in bold.

Property	Goal	No. 1	No. 2	No. 3	HSX
No. of coils per hfp	↓	**4**	5	6	6
$\langle \boldsymbol{B}\cdot\boldsymbol{n}\rangle/\langle B\rangle \times 10^{-3}$	↓	11	7.5	**5.3**	5.8
Max $(\boldsymbol{B}\cdot\boldsymbol{n}/B)\times 10^{-2}$	↓	4.7	4.2	**2.2**	2.5
Max Force [kN/m]	↓	0.7	0.88	0.7	**0.6**
Coil Length per hfp [m]	↓	**10.0**	12.1	13.4	13.4
Min CC Distance [m]	↑	**0.10**	**0.10**	0.09	0.09
Min CS Distance [m]	↑	**0.15**	0.13	0.11	0.14
Max κ [m^{-1}]	↓	9.7	12	**8**	12.34
Max MSC [m^{-1}]	↓	**25**	30	30.0	45.1
Quasisym. Error $\times 10^0$	↓	40.6	**37.6**	39.2	**37.6**
HTS length [km]	↓	24	23.2	**21.4**	**21.4**

E. Helically symmetric experiment (hsx)

The Helically Symmetric eXperiment (HSX) is a stellarator located at the University of Wisconsin in Madison. The device began operation in 1999 and is still the only device in the world with a quasihelical stellarator field. It is a four-field-period stellarator with 48 copper coils, around 1 T on axis, a major radius of 1.2 m, and an aspect ratio of 10. HSX was designed to demonstrate reduced direct orbit losses, Mercier stability, and reduced neoclassical transport. Due to its rotational transform of 1.05 on axis and quasihelical field, its neoclassical confinement is similar to a $q = 1/3$ tokamak [6]. It has also managed to demonstrate reduced flow damping and reduced electron heat transport [69,70]. Unlike the QH stellarator equilibrium shown in Sec. III B, which was optimized for low quasisymmetric error throughout the volume, HSX was only optimized for quasisymmetry on the last closed flux surface because at the time no robust optimization methods had been developed yet. Similarly to W7-X, HSX coils were obtained using a winding surface method, which yielded six coils per half-field period with high curvatures. From the twisted filaments, a copper winding pack was constructed for the final device; these are shown in Fig. 7. The procedure in this section does not aim at improving the forces, since the force requirements become much more stringent at reactor scale. Given the smaller dimensions and relatively low magnetic fields of this device, we focus instead on finding simpler coils that are easier to manufacture. However, HSX coils are made of copper, which means that unlike HTS coils curvature does not pose a risk to the integrity of the materials. We propose three different configurations with four, five, and six coils per half-field period, summarized in Table VI. Note that the target configuration used in this section is a free-boundary surface, meaning that it already has coil ripple embedded into it, making it harder to optimize for different numbers of coils per hfp and partially responsible for the high quasisymmetry errors observed in Table VI. The reference coilset used here is obtained by taking the center

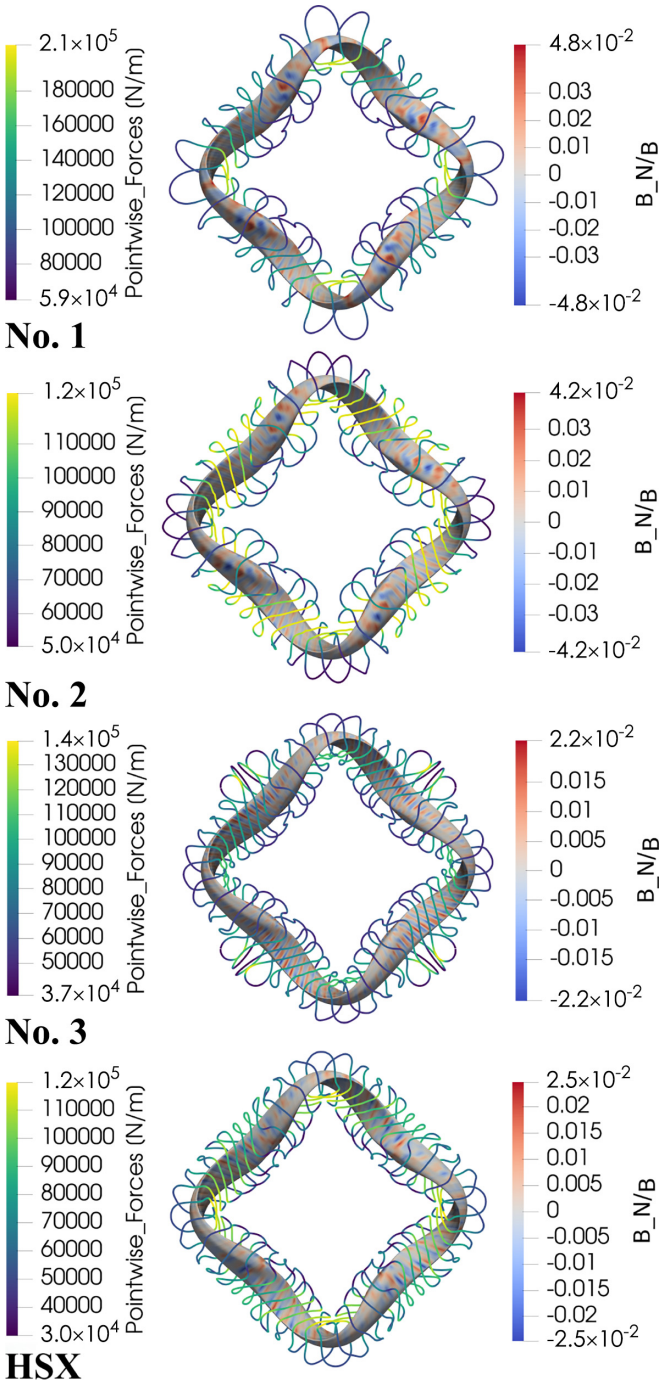


FIG. 7. New coil configurations for the HSX reactor optimized with the augmented Lagrangian method. (No. 1) First alternative configuration for the HSX stellarator with four coils per half-field period, simpler coils, and improved coil-to-surface and intercoil distance. (No. 2) Second alternative coilset for the HSX stellarator with five coils per half-field period and simpler coils. (No. 3) Third alternative coilset for HSX with improved field error, similar coil length, and considerably simpler coils. (No. 4) Original HSX coil design with six coils per half-field period. Forces appear very large because they are reported here before dividing by the number of turns.

spline of finite-built coils, which accounts for the field error being nonzero. Nevertheless, we report comparably low field errors, and as expected, the average field error progressively

increases with decreasing number of coils. This is consistent with the results from the previous QH, Stellaris, and W7-X configurations.

The three proposed configurations are mostly characterized by their reduction in curvature, which can be understood as a proxy for coil complexity. The reduction in mean-squared curvature reaches 45% in configuration No. 1, and curvature is decreased by 35% in configuration No. 3. Despite configuration No. 3 exhibiting the same number of coils per hfp, it manages to improve on the field accuracy while having simpler coils, the same clearance between coils, and losing only 3 cm in the coil-to-surface distance. The forces are generally increased in all the optimized configurations, as they are not a target. It remains the sole aspect where the original HSX design is better.

Configuration No. 1 is an exemplary solution. Configuration No. 1 has reduced by two the number of coils per half field period. This level of coil reduction has not been tried in previous sections, as it is often expected that field accuracy will not be precise enough with so few coils per hfp. Amazingly, and despite the fewer coils per hfp, configuration No. 1 has reasonable field accuracy, and, compared with the original six-coil design, improved minimum coil-coil and coil-surface distances, lower curvatures, and achieves all of these improvements with 25% less total coil length. Despite the factor of 2 increase in field error compared with the original design, the magnetic-field contours plotted in Fig. 13 still display accuracy comparable to the other candidates. As previously mentioned, fewer coils are of particular interest in a research-scale reactor where additional diagnostics are usually designed throughout the operation time and need space for new ports. The additional space in between coils on the outboard side is noticeable in Fig. 7. Fewer coils also means less machining time and workforce required in the case of HSX, as the coils are made out of copper and usually require the operation of CNC milling or other heavy machinery devices. This can directly translate to reduced costs, which is often a limiting factor in university-scale reactors and more generally in any type of reactor.

IV. CONCLUSION

In this work, we have introduced and validated an optimization methodology that fundamentally advances the state-of-the-art in stellarator coil design. The challenge of designing magnetic coils for a fusion reactor has long been characterized by difficult trade-offs between achieving the precise magnetic-field geometry required for plasma confinement and satisfying practical engineering constraints. Existing methods often lead to computationally intensive searches through vast parameter spaces, with no guarantee of arriving at a viable, let alone optimal, solution. Our approach confronts this challenge, establishing a robust and computationally efficient pathway for generating coil sets that are not only consistent with desired plasma equilibria but are also potentially manufacturable and meet the stringent requirements for reactor integration. Our methodology was successfully implemented and tested using the SIMSOPT framework.

By applying an augmented Lagrangian approach to coil optimization, it was possible to address three critical issues: (1) the automated adjustment of weights for satisfying the constraints in a nonconvex optimization problem, (2) the need for large computational resources, and (3) optimally satisfying the trade-off between low field errors and engineering constraints. To showcase the high performance of our method, five very different stellarator equilibria were considered, including quasi-axisymmetric, quasihelically symmetric, and quasi-isodynamic designs. New alternative coil sets were found for two operating stellarators, Wendelstein 7-X and HSX. This demonstrates the versatility of this method in order to generate cutting-edge coils independently of the plasma boundary. The quasi-axisymmetric design illustrated one of the main points of this work: the ability to perform single-optimization runs that achieve known Pareto-optimal solutions. By relaxing various constraints of the coil configurations, we were able to access new Pareto boundaries. The automated weight tuning and reliance on upper and lower bounds allow one to specify loose parameter space bounds, increasing the space of possible solutions.

All optimizations were performed on an Intel core i7–12650H laptop on four cores for Python multithreading and on one core for the evaluations of the functions and gradients. Future work includes accelerating the design of stellarator fusion reactors using modular coils, application to dipole coil arrays, incorporation into single-stage optimization efforts, and further algorithmic improvements.

ACKNOWLEDGMENTS

This work was supported through grants from the Simons Foundation under Award No. 560651 and from the Helmholtz Association Young Investigators Group program as Project No. VH-NG-1430. We thank Jorrit Lion, Gabriel Plunk, Alan Goodman, Matt Landreman, and Benedikt Geiger for sharing plasma boundary configuration files, and Rory Conlin, Paul Huslage, David Bindel, and Tom Simpson for useful discussions.

DATA AVAILABILITY

The configuration files to generate the figures and results in the present work can be found online [39].

APPENDIX A: ORDER INVESTIGATION

This section provides a brief investigation into the specific effects of varying order M on the quality of the optimized magnetic field via the Augmented Lagrangian method. For simplicity, we focus on the Landreman-Paul QA reactor-scale configuration. The quality of the coil solutions was quantified by $\langle \mathbf{B} \cdot \mathbf{n} \rangle / \langle \mathbf{B} \rangle$, as we found this was sufficient to draw conclusions. Orders 1 to 30 were investigated, and for each order, four optimization runs were performed with the Augmented Lagrangian method. The four optimizations were performed using four different initial conditions generated by adding Gaussian perturbations ($\sigma = 0.01$ and $L = 0.5$) to the initial set of coils. The results are summarized in Fig. 8.

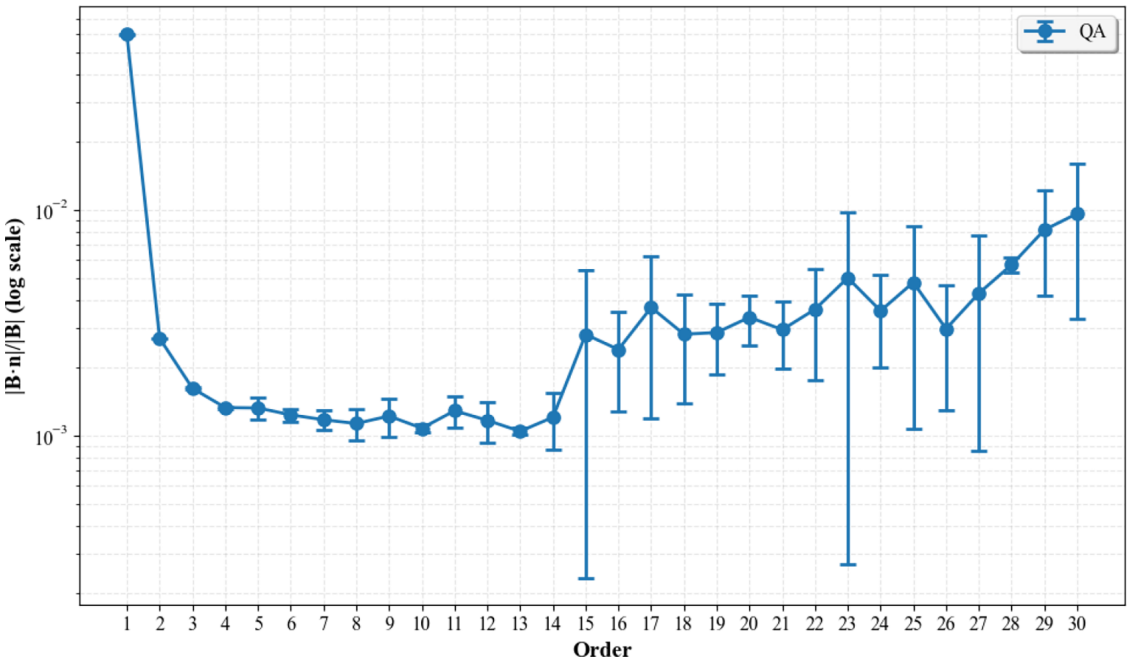


FIG. 8. Log-scale error bar plot for normal field error vs order for the QA configuration after ALM optimization.

Surprisingly, $M = 2$ is already sufficient to get most of the benefits of higher order, although orders 8–14 provide the most consistently low $\langle \mathbf{B} \cdot \mathbf{n} \rangle / \langle \mathbf{B} \rangle$ errors without adding overwhelming complexity. As $M \geqslant 15$ or so, the normalized field error begins to worsen (on average) and increase in variation from run to run, despite the extra degrees of freedom in the problem. Since the variation between runs is high, the problem may be increasingly nonconvex as M increases. This result is consistent with the traditional explanation that high Fourier modes are complicating the optimization landscape by introducing many suboptimal local minima [71].

Lastly, we found a similar trend using traditional coil optimization techniques and with different stellarators (although the optimal M choice can depend on the stellarator). In fact, with a traditional method, it was easy to generate final solutions with infeasible coil sets, where the coils were interlinking or even intersecting with the plasma surface, if the weights were not carefully tuned.

Initially, we found similar issues with the augmented Lagrangian approach. A critical consideration was that if the initial weight on the coil-surface distance was too small, the first subproblem solved by augmented Lagrangian could generate a solution with coil-surface intersection, after which all subsequent optimizations struggled to recover from. This first subproblem issue can be permanently solved by starting the optimization with an aggressive initial weight on the coil-surface distance penalty, after which the weight is tuned by the algorithm as before. We then manually checked in Fig. 8 that all of the resulting coil optimizations were feasible. In other words, we were able to completely remove infeasible coil solutions from the scan, and yet the trend with M remains. In summary, this experiment generally suggests that orders in the 6–15 range provide more consistent results for minimizing the normalized field error. Moreover, it suggests that if high-order coils are desired, a Fourier continuation approach is likely to perform better than direct optimization starting with a large value of M. As far as we are aware, this is standard practice in stage-1 but not stage-2 optimizations, so continuation should be carefully considered by future researchers.

APPENDIX B: QUASISYMMETRY METRIC

In this work, the metric used to quantify the departure from quasisymmetry is the two-term metric [9,72], which does not require a transformation into Boozer coordinates and can be expressed as

$$f_{QS} = \sum_{s_i} w_i \left\langle \left(\frac{1}{B^3} [(N - \iota M) \boldsymbol{B} \times \nabla B \cdot \nabla \psi - (MG + NI) \boldsymbol{B} \cdot \nabla B] \right)^2 \right\rangle. \tag{B1}$$

Here it appears as a weighted sum across all specified flux surfaces s_i, with the weights w_i by default set to 1, B is the magnetic-field amplitude, (N, M) are parameters specifying the class of quasisymmetry (either QH or QA), ψ is the toroidal flux, ι is the rotational transform, and G is the poloidal and I is the toroidal currents inside the surface.

APPENDIX C: EQUILIBRIA VERIFICATION

This section provides postprocessing illustrations demonstrating that the physics properties of the plasma are well-conserved for the coil solutions identified in this work. Figure 9 illustrates Boozer and Poincaré plots for a Landreman-Paul QA coil solution, and similarly for Landreman-Paul QH in Fig. 10, Stellaris in Fig. 11, W7-X in Fig. 12, and HSX in Fig. 13.

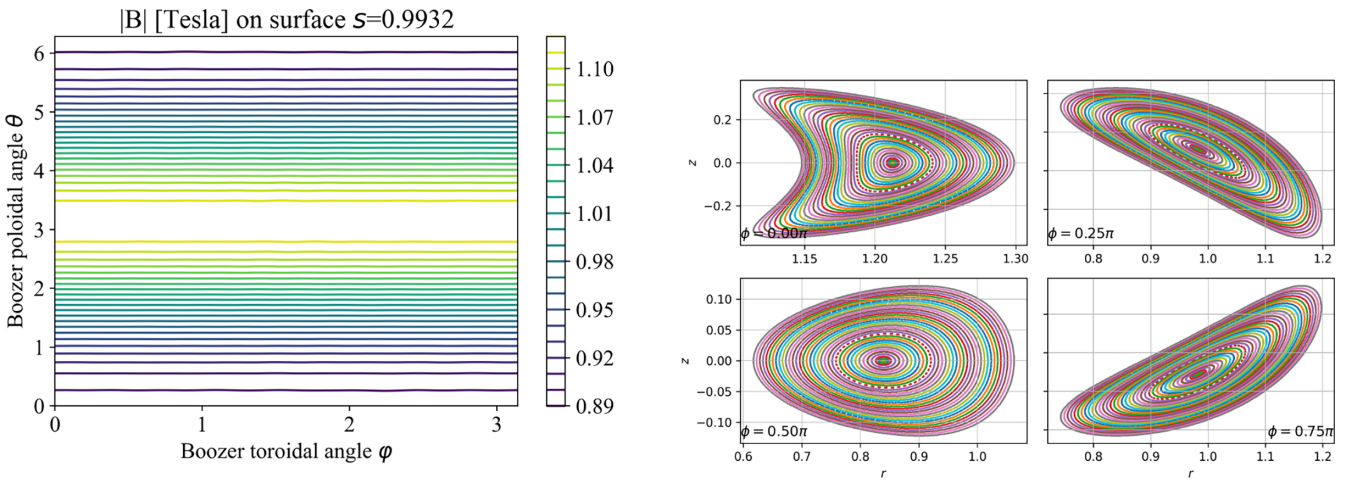


FIG. 9. (Left) Boozer plot of one of the ALM 25 m coil lengths from Fig. 2 configuration showing the quality of the quasisymmetry of the last closed flux surface, which delimits the plasma boundary. Perfectly straight lines mean that a very precise level of quasisymmetry is achieved. (Right) Poincaré plot showing a cross-section of the plasma where multiple nested flux surfaces are visible with no island chains.

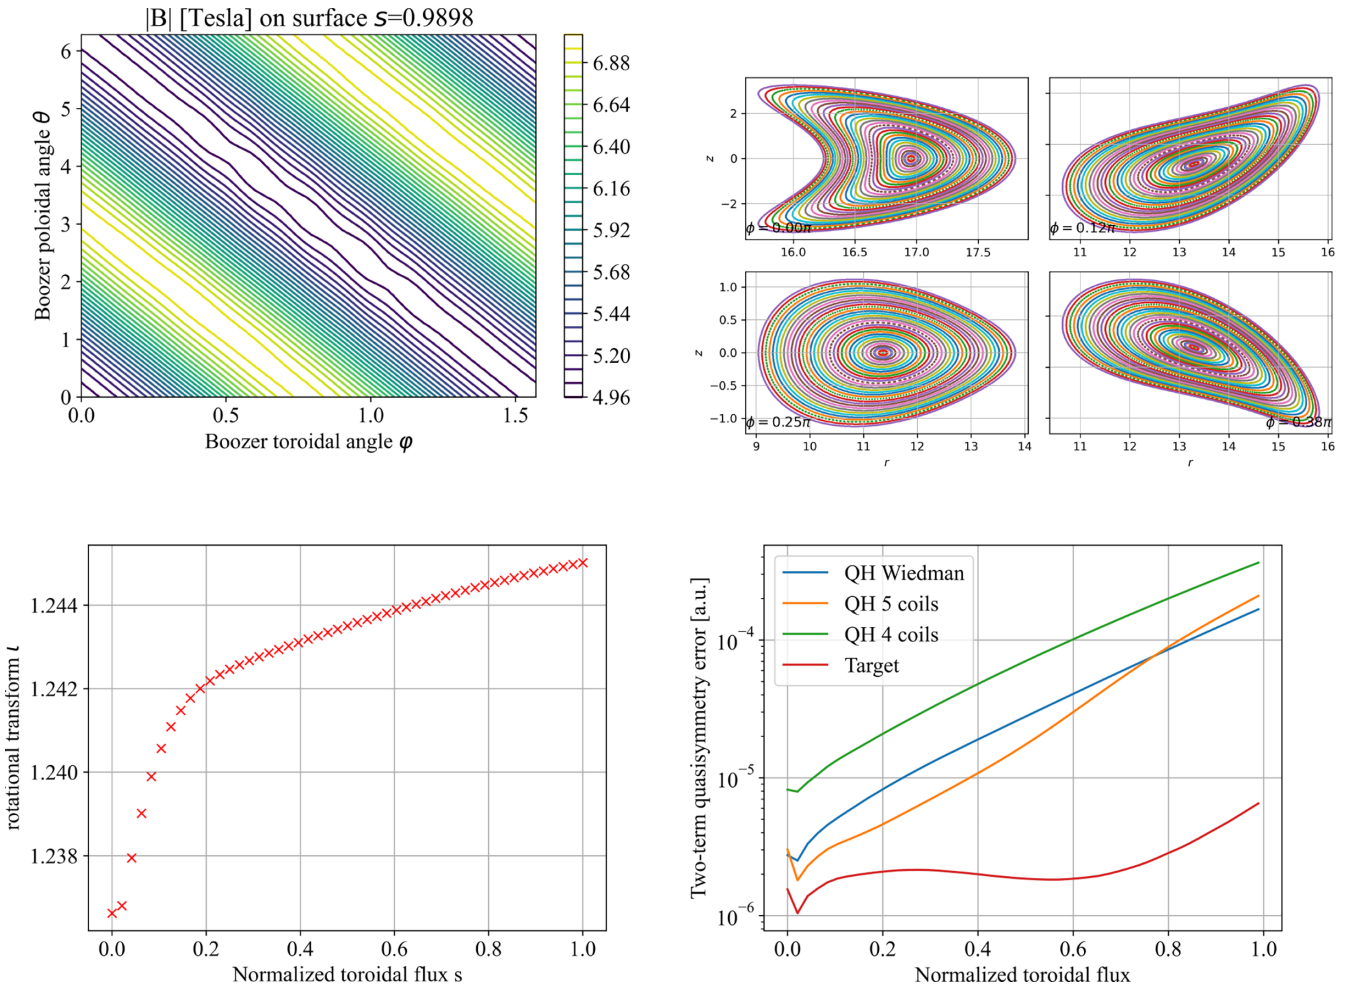


FIG. 10. (Top Left) Boozer plot for configuration No. 1 showing the quality of the quasihelical symmetry of the last closed flux surface. Perfectly straight lines mean that a very precise level of quasisymmetry is achieved. (Top Right) Poincaré plot of configuration No. 1 showing a cross section of the plasma where multiple nested flux surfaces are visible with no island chains and a rotational transform profile of the configuration which matches the ideal profile. (Bottom Left) Rotational transform profile of configuration No. 1. (Bottom Right) QS profiles for the configurations No. 1, No. 2, Wiedman, and the original QH target plasma.

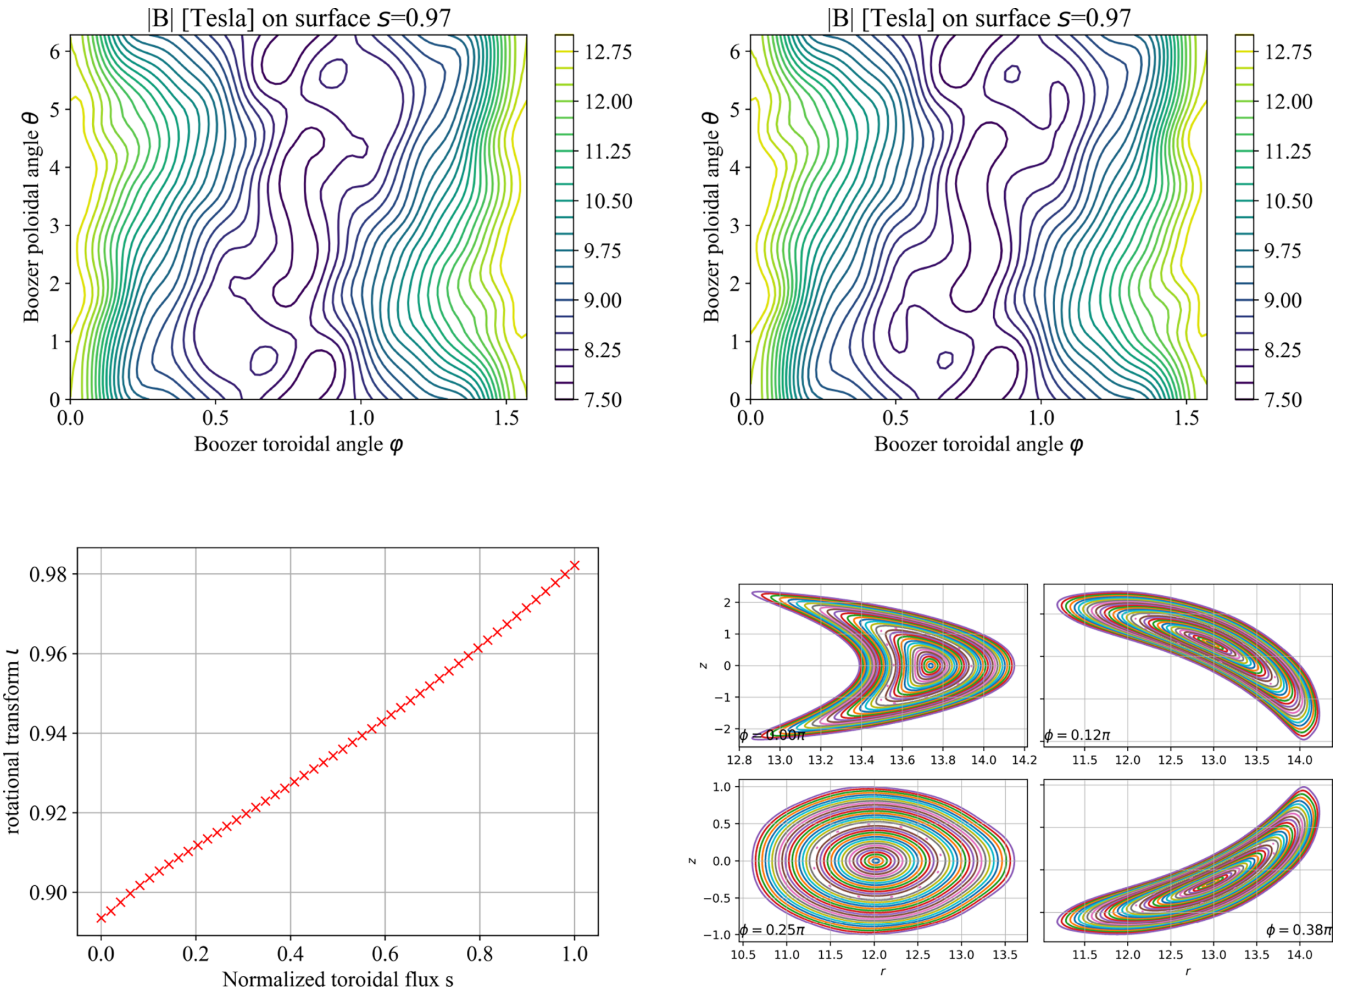


FIG. 11. (Top Left) Boozer plot of the No. 2 Stellaris configuration showing the quality of the quasi-isodynamic field generated by the coils (Top Right) Boozer plot of the target isodynamic field of Stellaris. The coils appear to accurately capture the desired features of the boundary magnetic surface. (Bottom Left) Rotational transform profile of the No. 2 Stellaris configuration. (Bottom Right) Poincaré plots of the No. 2 configuration.

APPENDIX D: FAST PARTICLE SIMULATIONS

Additional checks are performed regarding the confinement of α particles. This is performed for the Landreman-Paul QA three-coils-per-hfp solution versus the Wechsung L18 with four coils per hfp configuration in Fig. 14. Moreover, our Stellaris configurations No. 2 and No. 3 are also compared with the original Stellaris coil set in Fig. 15. The results here displayed correspond to the results of Ref. [38]. The simulations are performed with SIMPLE [73,74], a simplectic integrator to calculate fast ion orbits and losses. The simulations are performed by initializing 1000 particles with random pitch angles at 3.5 MeV at the radial coordinate of 0.2 and are followed for 0.2 s.

While configurations No. 2 and Stellaris display no losses after 0.2 seconds, configuration No. 3 performs worse with 3.5% losses. This is expected as this configuration is intentionally optimized to present simpler coils and therefore a worse field accuracy. Nonetheless, the losses remain very low and comparable to other top candidates for reactor designs self-consistently modeled with a coil set.

APPENDIX E: DIPOLE OPTIMIZATION

To illustrate that our method can easily integrate with any existing functionality in the coil optimization literature, we now present dipole optimizations using the Schuett-Henneberg (a tokamak-stellarator hybrid) QA equilibrium [75] as target field following a similar setup as Ref. [48]. Two optimizations are performed, one with the augmented Lagrangian Method and one with a standard stage two script. Both use 16 planar coils per half-field period as field-shaping coils. Additionally, two nonplanar toroidal field coils per half-field period are optimized. The toroidal field and dipole coils are jointly optimized, and the dipole coils can move in space and change current but are identical fixed-shape circular coils. The optimization targets and results used are presented in Table VII. The resulting configurations are shown in Fig. 16.

Note that the standard stage II optimization is not performed with random weights, but with already chosen weights to give better performance following the script used in Ref. [48]; however, without force and torque optimizations for the coils. The sole modification was to replace the minimizer of

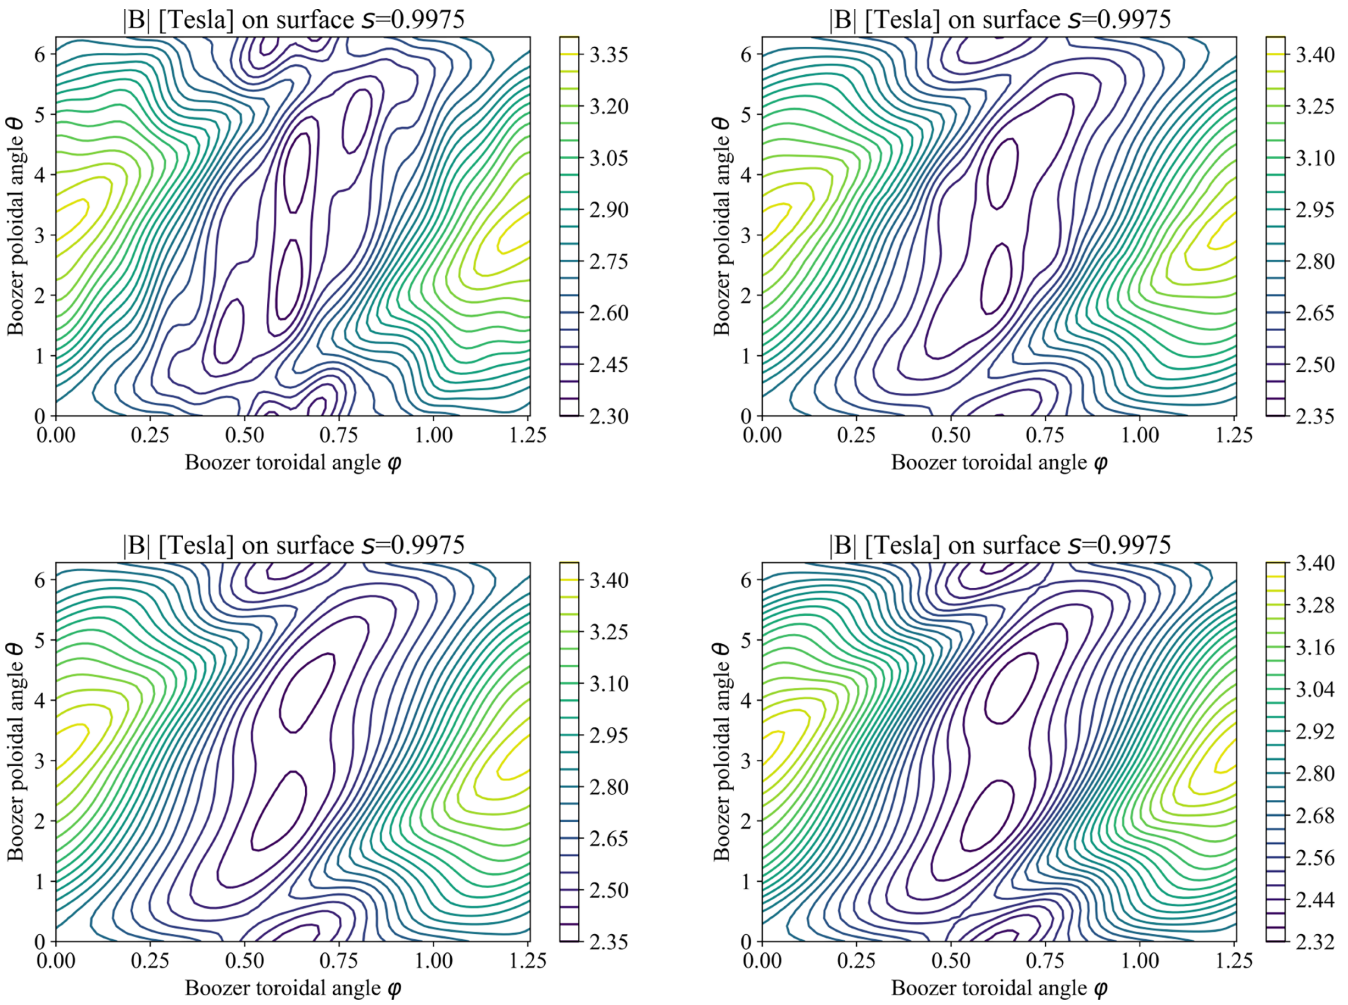


FIG. 12. Boozer plots of the four W7-X configurations presented in this work showing the quality of the omnigenous field generated by the coils: four coils per half-field period (No. 1 Top Left), five coils per hfp with 45 m coil length and reduced forces (No. 2 Top Right), five coils per hfp with 43 m coil length and best field accuracy (No. 3 Bottom Left) and the target W7-X field in fixed-boundary (Bottom Right).

the original script with the augmented Lagrangian method, and it performed as expected: apart from the coil-to-coil and coil-to-surface distances, which remain within 10% tolerance of the target field, the augmented Lagrangian outperforms the standard stage II optimization in most of the metrics. Notably, it reaches lower field errors with curves that exhibit much smaller maximum curvature and even 50 cm less conductor length.

APPENDIX F: COMPARISON WITH QUADRATIC PENALTY METHODS

To benchmark our work against standard penalty methods, we decided to perform a comparison between the augmented Lagrangian algorithm in question and the Quadratic Penalty Method (QPM) by Nocedal & Wright Framework 17.1 [45]:

For each outer iteration, the inner loop minimizes the augmented objective

$$Q(x,\mu) = f(x) + \frac{1}{2}\sum_i \mu_i c_i(x)^2 \tag{F1}$$

with L-BFGS-B to a tolerance τ_κ and between outer iterations the penalty grows: if the previous inner solve converged in at most 1000 iterations, $\mu_{\kappa+1} = 10\mu_\kappa$; otherwise $\mu_{\kappa+1} = 2\mu_\kappa$, τ_κ is halved each outer step toward 10^{-8}, this follows strictly the implementation presented in Framework 17.1. All the solves had 100 outer loop iterations, ensuring that both algorithms converged.

The constraints are the same ones as used previously in this work, with the addition of the arclength penalty that penalizes the quadrature point arclength variation.

We therefore present the results of the two mentioned coil optimization algorithms applied to the same stellarator coil-design problem, with exactly the same constraints, run on two target surfaces (for a QA and a QI stellarator) and a sweep of total coil-length thresholds. Both algorithms start from the same physical coilset (the same build problem instance, with coils placed along the magnetic axis for the QI surface) and use identical initial penalty parameters ($\mu_0 = 10$ broadcast scalar), so any difference observed is a consequence of the update rule, not of the starting point.

As a filtering policy, we only include here configurations that did not exhibit an individual constraint violation above

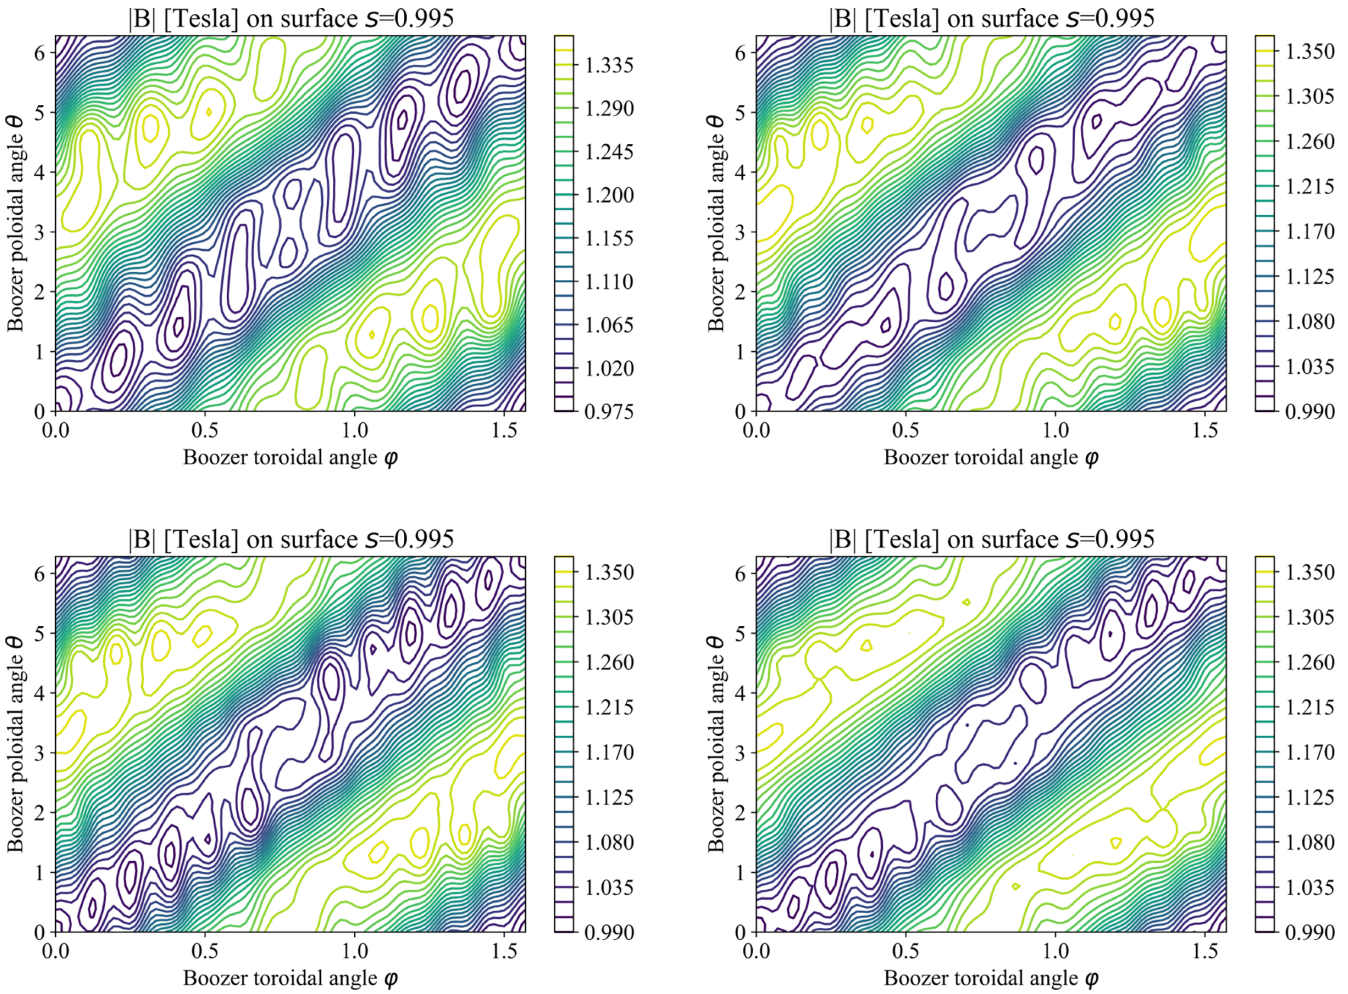


FIG. 13. Boozer plots of the four HSX configurations presented in this work showing the quality of the quasihelically symmetric field generated by the coils: four coils per half-field period (No. 1 Top Left), five coils per hfp (No. 2 Top Right), six coils per hfp (No. 3 Bottom Left) and the original HSX coil set (Bottom Right).

20% in order to be able to perform the comparison. As a figure of merit, we compare the maximum force on the coil versus the normalized squared flux, just like in the submitted paper.

Each data point corresponds to one coil set from a full ALM or QPM run; the thresholds are fixed during each optimization. Increasing the coil length upper threshold yields the different data points for both the ALM and QPM. The

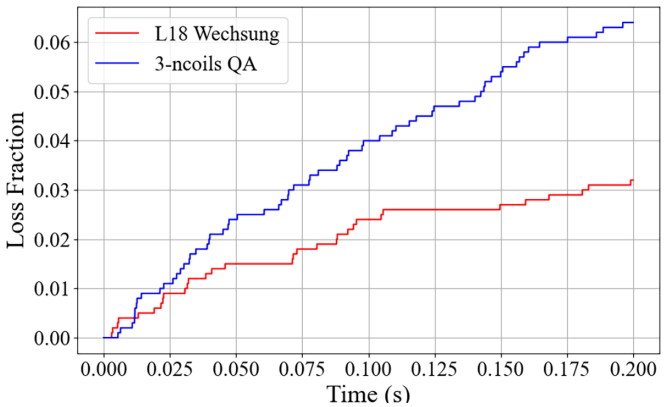


FIG. 14. Loss fractions for the Landreman-Paul QA three coils per hfp solution versus the Wechsung L18 with four coils per hfp after 0.2 s.

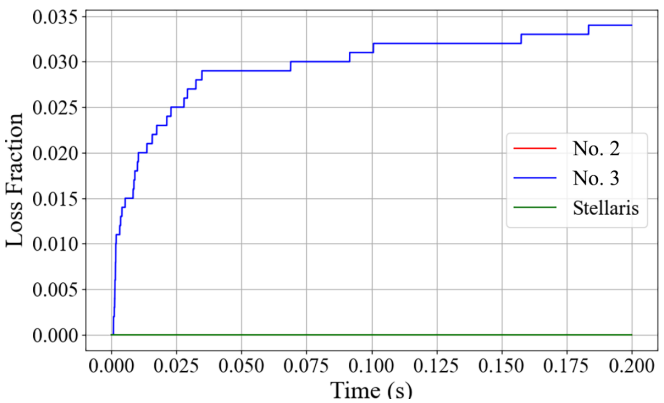


FIG. 15. Loss fractions for configurations No. 2, No. 3, and the original Stellaris coilset after 0.2 s. Note that both the red and the green lines have losses of 0%.

TABLE VII. Comparison of magnetic-field and coil properties for the two-field-period Schuett-Henneberg [1] QA equilibrium. Apart from the coil-to-coil and coil-to-surface distances, which remain within 10% tolerance of the target field, the augmented Lagrangian outperforms the standard stage II optimization in most of the metrics. Top-performing metrics are highlighted in bold. TF stands for Toroidal Field.

Property	Aug. Lag.	Std.	Thresh.
No. of coils per hfp ↓	18	18	
$\langle \boldsymbol{B}\cdot\boldsymbol{n}\rangle/\langle B\rangle \times 10^{-3}$ ↓	**1.9**	2.6	0
Max $(\boldsymbol{B}\cdot\boldsymbol{n}/B)\times 10^{-3}$ ↓	**9.6**	13	0
TF Coil Length per hfp [m] ↓	**80.04**	80.47	80
Min CC Dist. (All coils) [m] ↑	0.46	**0.5**	0.5
Min CC Dist. (TF coils) [m] ↑	0.93	**1**	1
Min CS Dist. [m] ↑	1.43	**1.48**	1.5
Max κ [m^{-1}] ↓	**0.59**	1.68	0.5
Max MSC [m^{-1}] ↓	**0.13**	0.14	0.05

arrows indicate how the data points move by increasing the upper length threshold for the ALM (red) and the QPM (blue), respectively. The runs were done for a maximum [15, 17.5, 20, 22.5, 25] m total coil length per hfp for the QA equilibrium and [12, 13, 14, 15, 16] m for the QI equilibrium, both with four coils per half-field period. So, the first data points of the QA ALM and QPM series were optimized for a maximum coil length of 15 m, the second for 17.5 m, and so on.

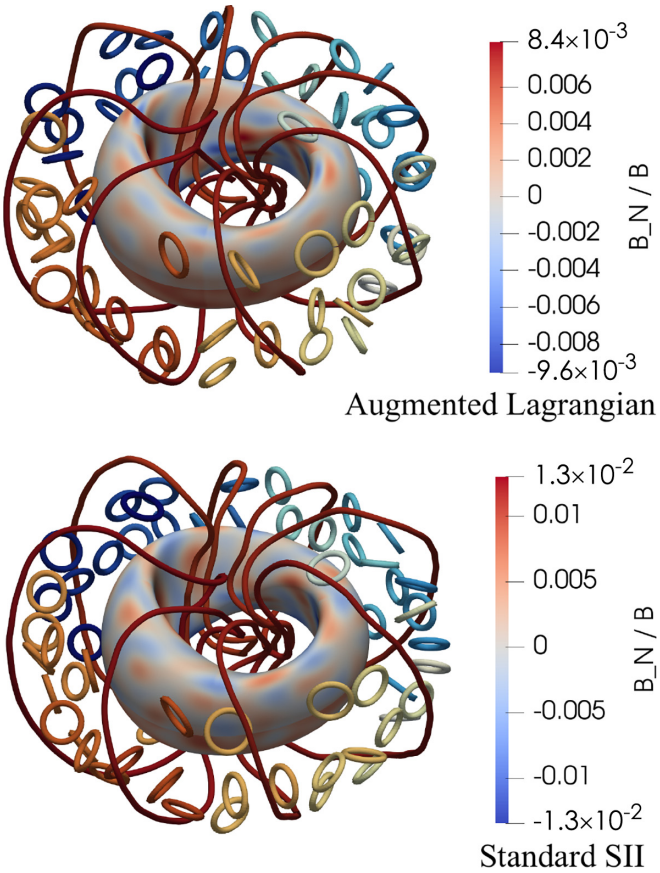


FIG. 16. Coils and target surfaces for the Schuett-Henneberg tokamak-stellarator hybrid.

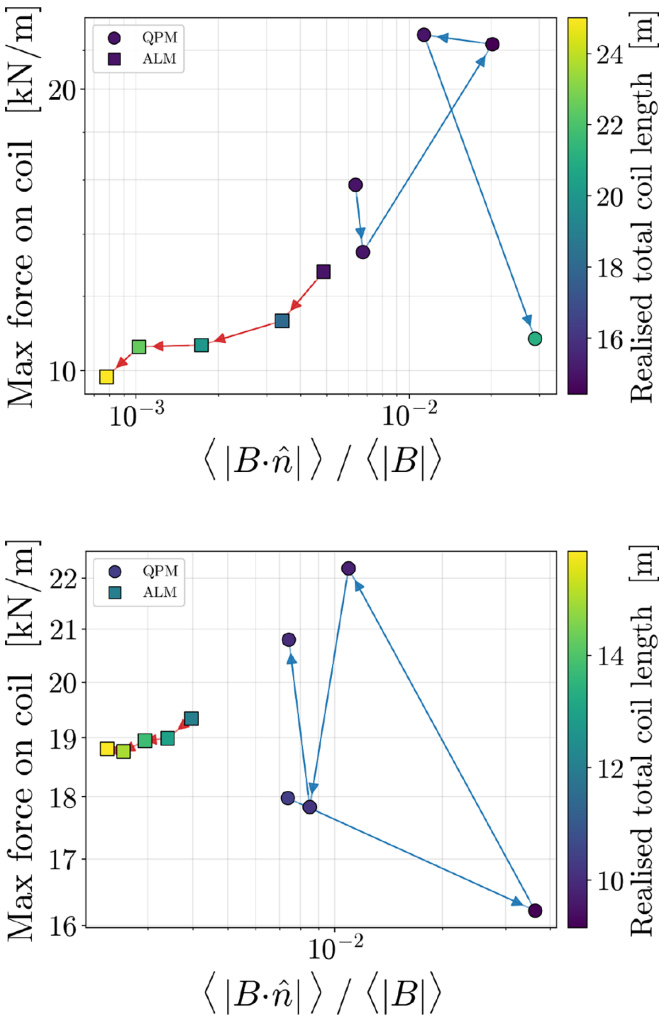


FIG. 17. (Top) Comparison between the Quadratic Penalty Method (QPM) and the Augmented Lagrangian Method (ALM); shown is the maximum force on coil versus normalized squared flux for varying coil-length thresholds for the Landreman-Paul QA configuration [9]. (Bottom) Same plot, but for completeness, the optimizations are performed on a Quasi-Isodynamic configuration [30]. Note that the colorbar corresponds to the realized total coil lengths of the coilsets.

The colorbar indicates the actual realized total coil length per half-field period in the final optimized coilsets.

The augmented Lagrangian (red arrows) appears in both cases to improve consistently both in squared flux and in maximum force as the length constraint is relaxed. The QPM trajectory, however, appears to be erratic as the length threshold is relaxed. As a reminder, the arrows in Fig. 17 show how the algorithm behaves by progressively relaxing the coil length constraint, i.e., by allowing the coils to get bigger. The irregular behavior of the QPM trajectory can be explained by observing that the realized coil length for the QPM appears to stay stuck even as the coil length threshold is relaxed. Moreover, the ALM pushes the coils to match the exact coil length threshold each time, as they are tight constraints. This is a revealing difference in behavior, and these findings strongly suggest the ALM algorithm

regularly outperforms QPM on this class of optimization problems.

One possible explanation behind the QPM getting stuck is that it does not possess the additional decision node regarding the constraint violation that the ALM has, which comes after one minimizer solve. This results in always increasing magnitudes of the penalty parameters μ_p of the violated constraints, leaving the unviolated constraints, like the coil length, untouched. However, in the ALM, depending on a true or false constraint violation up to a certain threshold η_p, it has the additional freedom of updating both the Lagrange multipliers and tightening the constraint violation threshold η_p while keeping the penalty parameters μ_k untouched. The QPM algorithm does not possess adaptive constraint violation thresholds.

Lastly, we make no claim that ALM always outperforms QPM. This depends on the nonconvexity and structure of the problem, and we have only shown evidence that on typical apples-to-apples problems in stellarator coil optimization, ALM appears to outperform QPM.

[1] F. C. Schuller, Disruptions in tokamaks, Plasma Phys. Control. Fusion **37**, A135 (1995).

[2] G. Jahns, M. Soler, B. Waddell, J. Callen, and H. Hicks, Internal disruptions in tokamaks, Nucl. Fusion **18**, 609 (1978).

[3] P. Helander, C. D. Beidler, T. M. Bird, M. Drevlak, Y. Feng, R. Hatzky, F. Jenko, R. Kleiber, J. H. E. Proll, Y. Turkin, and P. Xanthopoulos, Stellarator and tokamak plasmas: A comparison, Plasma Phys. Control. Fusion **54**, 124009 (2012).

[4] G. Grieger, W. Lotz, P. Merkel, J. Nührenberg, J. Sapper, E. Strumberger, H. Wobig, R. Burhenn, V. Erckmann, U. Gasparino, L. Giannone, H. J. Hartfuss, R. Jaenicke, G. Kühner, H. Ringler, A. Weller, F. Wagner, the W7-X Team, and the W7-AS Team, Physics optimization of stellarators, Phys. Fluids B **4**, 2081 (1992).

[5] T. Klinger, C. Baylard, C. Beidler, J. Boscary, H. Bosch, A. Dinklage, D. Hartmann, P. Helander, H. Maßberg, A. Peacock, T. Pedersen, T. Rummel, F. Schauer, L. Wegener, and R. Wolf, 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 Eng. Des. **88**, 461 (2013).

[6] F. Anderson, A. Almagri, D. Anderson, P. Matthews, J. Talmadge, and J. Shohet, Helically symmetric experiment, (HSX) goals, design and status, Fusion Technol. **27**, 273 (1995).

[7] S. P. Hirshman and J. C. Whitson, Steepest-descent moment method for three-dimensional magnetohydrodynamic equilibria, Phys. Fluids **26**, 3553 (1983).

[8] J. Nührenberg and R. Zille, Quasi-helically symmetric toroidal stellarators, Phys. Lett. A **129**, 113 (1988).

[9] M. Landreman and E. Paul, Magnetic fields with precise quasisymmetry for plasma confinement, Phys. Rev. Lett. **128**, 035001 (2022).

[10] D. W. Dudt and E. Kolemen, DESC: A stellarator equilibrium solver, Phys. Plasmas **27**, 102513 (2020).

[11] L.-M. Imbert-Gérard, E. J. Paul, and A. M. Wright, *An Introduction to Stellarators: From Magnetic Fields to Symmetries and Optimization* (Society for Industrial and Applied Mathematics, Philadelphia, 2024); https://epubs.siam.org/doi/pdf/10.1137/1.9781611978223.

[12] V. Erckmann, H.-J. Hartfuss, M. Kick, H. Renner, J. Sapper, F. Schauer, E. Speth, F. Wesner, F. Wagner, M. Wanner, A. Weller, and H. Wobig, The W7-X project: Scientific basis and technical realization, in *17th IEEE/NPSS Symposium Fusion Engineering (Cat. No.97CH36131)* (Institute of Electrical and Electronics Engineers (IEEE), Piscataway, NJ, 1997), Vol. 1, pp. 40–48.

[13] R. Strykowsky, T. Brown, J. Chrzanowski, M. Cole, P. Heitzenroeder, G. Neilson, D. Rej, and M. Viol, Engineering cost & schedule lessons learned on NCSX, in *2009 23rd IEEE/NPSS Symposium on Fusion Engineering* (IEEE (Institute of Electrical and Electronics Engineers), New York, NY, 2009), pp. 1–4.

[14] H.-S. Bosch, R. Brakel, T. Braeuer, V. Bykov, P. van Eeten, J.-H. Feist, F. Füllenbach, M. Gasparotto, H. Grote, T. Klinger, H. Laqua, M. Nagel, D. Naujoks, M. Otte, K. Risse, T. Rummel, J. Schacht, A. Spring, T. Sunn Pedersen, R. Vilbrandt, *et al.*, Final integration, commissioning and start of the Wendelstein 7-X stellarator operation, Nucl. Fusion **57**, 116015 (2017).

[15] F. Wechsung, A. Giuliani, M. Landreman, A. Cerfon, and G. Stadler, Single-stage gradient-based stellarator coil design: Stochastic optimization, Nucl. Fusion **62**, 076034 (2022).

[16] A. Giuliani, Direct stellarator coil design using global optimization: Application to a comprehensive exploration of quasi-axisymmetric devices, J. Plasma Phys. **90**, 905900303 (2024).

[17] R. Wu, T. Kruger, and C. Swanson, Planar coil optimization for the Eos stellarator using sparse regression, Plasma Phys. Control. Fusion **67**, 035019 (2025).

[18] R. Jorge, A. Giuliani, and J. Loizu, Simplified and flexible coils for stellarators using single-stage optimization, Phys. Plasmas **31**, 112501 (2024).

[19] Y. Suzuki, J. Huang, N. Wang, and Y. Ding, Design of simple stellarator using tilted toroidal field coils, Fusion Eng. Des. **173**, 112843 (2021).

[20] L. Fu, E. J. Paul, A. A. Kaptanoglu, and A. Bhattacharjee, Global stellarator coil optimization with quadratic constraints and objectives, Nucl. Fusion **65**, 026045 (2025).

[21] M. Drevlak, Optimization of heterogeneous magnet systems, *12th International Stellarator Workshop, Madison, WI* (1999).

[22] N. Pomphrey, L. Berry, A. Boozer, A. Brooks, R. Hatcher, S. Hirshman, L.-P. Ku, W. Miner, H. Mynick, W. Reiersen, D. Strickler, and P. Valanju, Innovations in compact stellarator coil design, Nucl. Fusion **41**, 339 (2001).

[23] D. J. Strickler, L. A. Berry, and S. P. Hirshman, Designing coils for compact stellarators, Fusion Sci. Technol. **41**, 107 (2002).

[24] T. Brown, J. Breslau, D. Gates, N. Pomphrey, and A. Zolfaghari, Engineering optimization of stellarator coils lead to improvements in device maintenance, in *2015 IEEE 26th Symposium on Fusion Engineering (SOFE)* (IEEE (Institute of Electrical and Electronics Engineers), Piscataway, NJ, 2015), pp. 1–6.

[25] J. Lion, J.-C. Anglès, L. Bonauer, A. Bañón Navarro, S. Cadena Ceron, R. Davies, M. Drevlak, N. Foppiani, J. Geiger, A.

Goodman, W. Guo, E. Guiraud, F. Hernández, S. Henneberg, R. Herrero, C. Hintze, H. Höchter, J. Jelonnek, F. Jenko, R. Jorge, *et al.*, Stellaris: A high-field quasi-isodynamic stellarator for a prototypical fusion power plant, Fusion Eng. Des. **214**, 114868 (2025).

[26] D. T. Anderson, J. M. Canik, C. C. Hegna, and C. M. Mowry, A comprehensive, unified baseline physics design for the type one energy stellarator fusion pilot power plant, 'Infinity Two', J. Plasma Phys. **91**, e65 (2025).

[27] D. Gates, S. Aslam, B. Berzin, P. Bonofiglo, A. Cote, D. Dudt, E. Flom, D. Fort, A. Koen, T. Kruger, S. Kumar, M. Martin, A. Ottaviano, S. Pasmann, P. Romano, C. Swanson, L. Tang, E. Winkler, and R. Wu, Stellarator fusion systems enabled by arrays of planar coils, Nucl. Fusion **65**, 026052 (2025).

[28] F. Volpe, Renaissance Fusion Technologies, https://renfusion.eu/technology (2023), accessed: 1 October 2023.

[29] J. Miyazawa and T. Goto, Development of steady-state fusion reactor by Helical Fusion, Phys. Plasmas **30**, 050601 (2023).

[30] A. Goodman, K. Camacho Mata, S. Henneberg, R. Jorge, M. Landreman, G. Plunk, H. Smith, R. Mackenbach, C. Beidler, and P. Helander, Constructing precisely quasi-isodynamic magnetic fields, J. Plasma Phys. **89**, 905890504 (2023).

[31] D. W. Dudt, A. G. Goodman, R. Conlin, D. Panici, and E. Kolemen, Magnetic fields with general omnigenity, J. Plasma Phys. **90**, 905900120 (2024).

[32] J. L. Velasco, I. Calvo, F. J. Escoto, E. Sánchez, H. Thienpondt, and F. I. Parra, Piecewise omnigenous stellarators, Phys. Rev. Lett. **133**, 185101 (2024).

[33] M. Landreman and P. J. Catto, Omnigenity as generalized quasisymmetry, Phys. Plasmas **19**, 056103 (2012).

[34] A. G. Goodman, P. Xanthopoulos, G. G. Plunk, H. Smith, C. Nührenberg, C. D. Beidler, S. A. Henneberg, G. Roberg-Clark, M. Drevlak, and P. Helander, Quasi-isodynamic stellarators with low turbulence as fusion reactor candidates, PRX Energy **3**, 023010 (2024).

[35] R. Nies, E. J. Paul, D. Panici, S. R. Hudson, and A. Bhattacharjee, Exploration of the parameter space of quasisymmetric stellarator vacuum fields through adjoint optimisation, J. Plasma Phys. **90**, 905900620 (2024).

[36] M. Landreman, Mapping the space of quasisymmetric stellarators using optimized near-axis expansion, J. Plasma Phys. **88**, 905880616 (2022).

[37] M. Landreman, B. Medasani, F. Wechsung, A. Giuliani, R. Jorge, and C. Zhu, SIMSOPT: A flexible framework for stellarator optimization, J. Open Source Softw. **6**, 3525 (2021).

[38] P. F. Gil, W. Li, J. Stratton, A. A. Kaptanoglu, and E. V. Stenson, companion paper, Stellarator coils for future fusion reactors via an augmented lagrangian approach, Phys. Rev. Lett. **137**, 065101 (2026).

[39] P. F. Gil, W. Li, J. Stratton, A. Kaptanoglu, and E. V. Stenson, Coilsets and Scripts from Augmented Lagrangian Methods for Stellarator Coils [Dataset], Zenodo, 2026, https://doi.org/10.5281/zenodo.18497939.

[40] P. Merkel, Solution of stellarator boundary value problems with external currents, Nucl. Fusion **27**, 867 (1987).

[41] M. Landreman, An improved current potential method for fast computation of stellarator coil shapes, Nucl. Fusion **57**, 046003 (2017).

[42] M. Landreman, B. Medasani, and C. Zhu, Stellarator optimization for good magnetic surfaces at the same time as quasisymmetry, Phys. Plasmas **28**, 092505 (2021).

[43] R. Conlin, P. Kim, D. W. Dudt, D. Panici, and E. Kolemen, Stellarator optimization with constraints, J. Plasma Phys. **90**, 905900501 (2024).

[44] A. R. Conn, N. I. M. Gould, and P. L. Toint, *Trust-Region Methods*, MPS/SIAM Series on Optimization (Society for Industrial and Applied Mathematics, Philadelphia, 2000).

[45] J. Nocedal and S. J. Wright, *Numerical Optimization*, 2nd ed. (Springer Science & Business Media, New York, NY, 2006).

[46] D. C. Liu and J. Nocedal, On the limited memory BFGS method for large scale optimization, Math. Program. **45**, 503 (1989).

[47] C. Zhu, S. R. Hudson, Y. Song, and Y. Wan, New method to design stellarator coils without the winding surface, Nucl. Fusion **58**, 016008 (2018).

[48] A. A. Kaptanoglu, A. Wiedman, J. Halpern, S. Hurwitz, E. J. Paul, and M. Landreman, Reactor-scale stellarators with force and torque minimized dipole coils, Nucl. Fusion **65**, 046029 (2025).

[49] S. Hurwitz, M. Landreman, and T. M. Antonsen, Efficient calculation of the self-magnetic field, self-force, and self-inductance for electromagnetic coils, IEEE Trans. Magn. **60**, 1 (2024).

[50] R. Robin and F. A. Volpe, Minimization of magnetic forces on stellarator coils, Nucl. Fusion **62**, 086041 (2022).

[51] S. Guinchard, S. R. Hudson, and E. J. Paul, Including the vacuum energy in stellarator coil design, Plasma Phys. Control. Fusion **67**, 035028 (2025).

[52] S. Hurwitz, M. Landreman, P. Huslage, and A. Kaptanoglu, Electromagnetic coil optimization for reduced Lorentz forces, Nucl. Fusion **65**, 056044 (2025).

[53] Z. S. Hartwig, R. F. Vieira, B. N. Sorbom, R. A. Badcock, M. Bajko, W. K. Beck, B. Castaldo, C. L. Craighill, M. Davies, J. Estrada, *et al.*, VIPER: An industrially scalable high-current high-temperature superconductor cable, Supercond. Sci. Technol. **33**, 11LT01 (2020).

[54] Z. Zhao, P. Moore, and L. Chiesa, Structural modeling of REBCO VIPER cable for high-field magnet applications, IEEE Trans. Appl. Supercond. **32**, 1 (2022).

[55] N. Riva, R. Granetz, R. Vieira, A. Hubbard, A. Pfeiffer, P. Harris, C. Chamberlain, Z. Hartwig, A. Watterson, D. Anderson, *et al.*, Development of the first non-planar REBCO stellarator coil using VIPER cable, Supercond. Sci. Technol. **36**, 105001 (2023).

[56] F. Wechsung, M. Landreman, A. Giuliani, A. Cerfon, and G. Stadler, Precise stellarator quasi-symmetry can be achieved with electromagnetic coils, Proc. Natl. Acad. Sci. USA **119**, e2202084119 (2022).

[57] F. Najmabadi, A. R. Raffray, S. I. Abdel-Khalik, L. Bromberg, L. Crosatti, L. El-Guebaly, P. R. Garabedian, A. A. Grossman, D. Henderson, A. Ibrahim, T. Ihli, T. B. Kaiser, B. Kiedrowski, L. P. Ku, J. F. Lyon, R. Maingi, S. Malang, C. Martin, T. K. Mau, B. Merrill, *et al.*, The ARIES-CS compact stellarator fusion power plant, Fusion Sci. Technol. **54**, 655 (2008).

[58] D. Spong, J. Harris, A. Ware, S. Hirshman, and L. Berry, Shear flow generation in stellarators—configurational variations, Nucl. Fusion **47**, 626 (2007).

[59] A. Wiedman, S. Buller, and M. Landreman, Coil optimization for quasi-helically symmetric stellarator configurations, J. Plasma Phys. **90**, 905900307 (2024).

[60] C. D. Beidler, H. M. Smith, A. Alonso, T. Andreeva, J. Baldzuhn, M. N. A. Beurskens, *et al.*, Demonstration of reduced neoclassical energy transport in Wendelstein 7-X, Nature (London) **596**, 221 (2021).

[61] P. Helander and J. Nührenberg, Bootstrap current and neoclassical transport in quasi-isodynamic stellarators, Plasma Phys. Control. Fusion **51**, 055004 (2009).

[62] P. Helander, Theory of plasma confinement in non-axisymmetric magnetic fields, Rep. Prog. Phys. **77**, 087001 (2014).

[63] H. E. Mynick, N. Pomphrey, and P. Xanthopoulos, Optimizing stellarators for turbulent transport, Phys. Rev. Lett. **105**, 095004 (2010).

[64] J. Proll, G. Plunk, B. Faber, T. Görler, P. Helander, I. McKinney, M. Pueschel, H. Smith, and P. Xanthopoulos, Turbulence mitigation in maximum-J stellarators with electron-density gradient, J. Plasma Phys. **88**, 905880112 (2022).

[65] P. Helander, J. Geiger, and H. Maaßberg, On the bootstrap current in stellarators and tokamaks, Phys. Plasmas **18**, 092505 (2011).

[66] C. Beidler, G. Grieger, F. Herrnegger, E. Harmeyer, J. Kisslinger, W. Lotz, H. Maassberg, P. Merkel, J. Nuhrenberg, F. Rau, *et al.*, Physics and engineering design for Wendelstein VII-X, Fusion Technology **17,** 148 (1990).

[67] CEA Irfu, Tests on the Wendelstein 7-X Stellarator Coils, Institut de Recherche sur les lois Fondamentales de l'Univers (Irfu), CEA Saclay, https://irfu.cea.fr/dacm/en/Phocea/Vie_des_labos/Ast/ast_visu.php?id_ast=3012.

[68] T. S. Pedersen, M. Otte, S. Lazerson, *et al.*, Confirmation of the topology of the Wendelstein 7-X magnetic field to better than 1:100,000, Nat. Commun. **7**, 13493 (2016).

[69] S. P. Gerhardt, J. N. Talmadge, J. M. Canik, and D. T. Anderson, Measurements and modeling of plasma flow damping in the Helically Symmetric eXperiment, Phys. Plasmas **12**, 056116 (2005).

[70] J. M. Canik, D. T. Anderson, F. S. B. Anderson, C. Clark, K. M. Likin, J. N. Talmadge, and K. Zhai, Reduced particle and heat transport with quasisymmetry in the Helically Symmetric Experiment, Phys. Plasmas **14**, 056107 (2007).

[71] B. Jang, R. Conlin, and M. Landreman, Exponential spectral scaling: Robust and efficient stellarator boundary optimization via mode-dependent scaling, arXiv:2509.16320.

[72] E. Rodríguez, E. Paul, and A. Bhattacharjee, Measures of quasisymmetry for stellarators, J. Plasma Phys. **88**, 905880109 (2022).

[73] C. G. Albert, S. V. Kasilov, and W. Kernbichler, Symplectic integration with non-canonical quadrature for guiding-center orbits in magnetic confinement devices, J. Comput. Phys. **403**, 109065 (2020).

[74] C. G. Albert, S. V. Kasilov, and W. Kernbichler, Accelerated methods for direct computation of fusion alpha particle losses within, stellarator optimization, J. Plasma Phys. **86**, 815860201 (2020).

[75] T. M. Schuett and S. A. Henneberg, Exploring novel compact quasi-axisymmetric stellarators, Phys. Rev. Res. **6**, L042052 (2024).