

Stochastic and a posteriori optimization to mitigate coil manufacturing errors in stellarator design

Florian Wechsung¹, Andrew Giuliani¹, Matt Landreman²,
 Antoine Cerfon¹, Georg Stadler¹

¹ Courant Institute of Mathematical Sciences, New York University, New York, USA

² Institute for Research in Electronics & Applied Physics, University of Maryland-College Park, Maryland, USA

E-mail: wechsung@nyu.edu, giuliani@cims.nyu.edu, mattland@umd.edu, cerfon@cims.nyu.edu, stadler@cims.nyu.edu

Abstract. It was recently shown in [Wechsung et. al., Proc. Natl. Acad. Sci. USA, 2022] that there exist electromagnetic coils that generate magnetic fields which are excellent approximations to quasi-symmetric fields and have very good particle confinement properties. Using a Gaussian process based model for coil perturbations, we investigate the impact of manufacturing errors on the performance of these coils. We show that even fairly small errors result in noticeable performance degradation. While stochastic optimization yields minor improvements, it is not able to mitigate these errors significantly. As an alternative to stochastic optimization, we then formulate a new optimization problem for computing optimal adjustments of the coil positions and currents without changing the shapes of the coil. These a-posteriori adjustments are able to reduce the impact of coil errors by an order of magnitude, providing a new perspective for dealing with manufacturing tolerances in stellarator design.

1. Introduction

One of the most expensive and technically challenging aspects of the construction of stellarators [14, 21, 24] is the design, manufacturing and assembly of the primary coil system that generates the magnetic field confining the plasma. As computing resources and simulation capabilities have increased over the last few decades, large-scale numerical optimization has proven to be the method of choice for the design of new stellarator experiments [11]. Numerical simulations of guiding-center trajectories have shown that the state-of-the-art stellarator designs obtained from these optimization efforts are now able to confine even highly energetic particles for long timescale with minimal losses — a key requirement for practical use as fusion reactors [4, 7, 16, 26]. However, good performance of a stellarator can only be guaranteed if the coils are manufactured and placed with high accuracy [1, 10, 23], a requirement that has turned out to be a significant cost driver and risk factor for any plans to build a stellarator experiment [21, 24]. To illustrate the fragility of confinement to even rather small coil errors, we consider coils [26] that approximate the two quasi-axisymmetric (QA) configurations found in [16]. We compute particle losses for alpha particles for the case when the coils are built perfectly, and when they have errors of a few millimeters, which may be compared to coil lengths of the order of a few meters (we give a precise

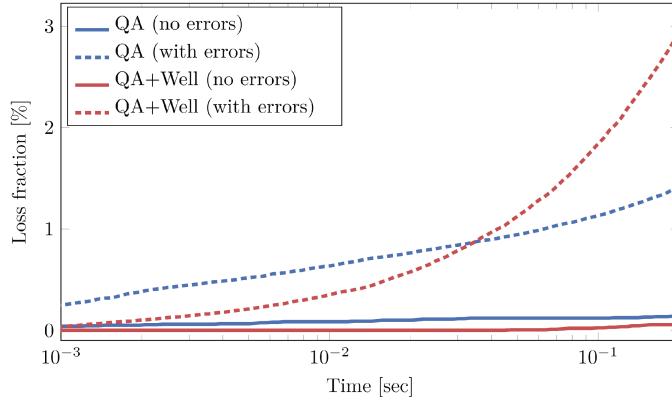


Figure 1. Loss fraction of alpha particles spawned at half-radius over time. Coil errors are approximately 1000 times smaller than the diameter of the device. We observe that such coil perturbations result in significant degradation of particle confinement.

description of the error model in Section 2), cf. Figure 1. We observe that particle losses have a significant relative increase in the presence of errors: from 0.14% to 1.56% for the design without a magnetic well, and from 0.06% to 3.05% for the design with a magnetic well. While these numbers remain small in absolute terms compared to previous stellarators ([16], Figure 6a), the increased losses are clearly undesirable, and may be larger than acceptable level for a steady-state fusion reactor [13, 28].

In an effort to lessen the impact of errors, stochastic optimization approaches have recently drawn significant attention in the stellarator community [9, 17–19, 27]. The core idea underlying the various approaches is the following: by including coil errors explicitly as part of the optimization problem, one hopes to find minimizers that perform well even when the coils are not built or placed perfectly. Using this approach, Lobsien et al. [19] were able to find coils for the W7-X experiment that on average improve their coil design objective by 20% compared to those obtained from standard deterministic optimization. Accordingly, a significant fraction of this article is focused on the merits and limitations of stochastic optimization for coil design.

While the methods we propose are generally applicable, in this work we focus on the vacuum magnetic fields found by Landreman & Paul [16]. These fields satisfy quasi-symmetry to a formerly unprecedented accuracy and have excellent confinement properties. Furthermore, it was recently shown in [26] that it is possible to design electromagnetic coils that reproduce these fields to high precision. However, as these calculations did not take coil errors into account, the goal of this present work is to understand, and then mitigate, the impact of such errors on the performance of the stellarator. This work contains two main contributions. First, we show that, perhaps surprisingly, stochastic optimization only yields marginal improvements, and we provide an explanation for these results in the context of our optimization problem. Second, we show that, on the other hand, an a-posteriori correction approach based on minor adjustments of the coils *after* they have been built is able to mitigate manufacturing errors to a significant degree. The idea behind our approach is similar to the coil correction process performed for the W7-X experiment during construction [3, 5], but differs in the choice of objective function and also considers

adjustments to the currents.

The manuscript is structured as follows: using the model for coil errors introduced in [27], we first perform a systematic study of the performance of the confining magnetic field as the magnitude of the error is increased. We find that the precise quasi-symmetry property is quickly lost even for small coil errors. In order to improve robustness, we then apply stochastic optimization to the problem. While this results in slightly smaller average field error, the effect is small and does not result in improved particle confinement. Motivated by this finding, we propose a novel strategy to deal with coil manufacturing errors. Assuming that the coils (and hence the errors that were made during production) can be measured accurately before placing them, we solve a separate optimization problem to find adjustments to the placement of the coils that mitigate manufacturing errors as much as possible. Assuming perfect measurements and placement, we show that these position adjustments are able to reduce the impact of manufacturing errors by several orders of magnitude. This result motivates a larger emphasis on the accurate measurement and placement of coils, at a level comparable to the attention given to their manufacturing.

2. Stochastic optimization for precise quasi-symmetry

2.1. Formulation of the coil optimization problem and model for coil errors

Following the approach of the FOCUS coil design tool [29], we solve an optimization problem to find coil shapes and currents so that the magnitude of the normal component of the magnetic field induced by the coils is minimized on the boundary S of a given ideal MHD equilibrium. We use the same setup as in [26]: for each of the two quasi axisymmetric fields of [16], 16 coils are optimized in order to reproduce the target field. Since the field satisfies two-fold rotational symmetry as well as stellarator symmetry, we only need to describe the shape and current of four independent modular coils to determine the whole system. We make the common single filament assumption and represent each coil as a smooth, periodic function $\mathbf{\Gamma}^{(i)} : [0, 2\pi] \rightarrow \mathbb{R}^3$, represented by a truncated Fourier series [29, eqn. (2)].

We consider an objective that is given by the relative quadratic flux on the target surface,

$$f(\mathbf{B}) = \int_S \left(\frac{\mathbf{B} \cdot \mathbf{n}}{\|\mathbf{B}\|} \right)^2 ds, \quad (1)$$

and denote the magnetic field induced by a set of coils by $\mathbf{B}(\mathbf{c}) = \mathbf{B}(\{\mathbf{\Gamma}^{(i)}(\mathbf{c})\}_{i=1}^{16}, \{I^{(i)}(\mathbf{c})\}_{i=1}^{16})$, i.e., the parameter vector \mathbf{c} contains the Fourier coefficients describing the four independent coils and their currents $I^{(i)}(\mathbf{c})$, and the remaining 12 coils and currents are determined through symmetries. We recall that if $f(\mathbf{B}) = 0$, then the field induced by the coils matches the target field exactly in the volume bounded by the surface.

Simply minimizing $\mathbf{c} \mapsto f(\mathbf{B}(\mathbf{c}))$ is an ill-posed problem in general, so in order to obtain coils that can be constructed at reasonable cost, the flux objective is usually combined with penalties or constraints on the geometry of the coils. In this work, we follow the formulation of [26] and consider inequality constraints on the gap between coils, their mean squared and maximum curvature, and their length. Shorter coils are usually less complex geometrically, which is why the coil length is a commonly used regularizer in coil optimization. In order to study the sensitivity of optima to this regularizer, we vary the coil length constraint, which bounds the sum of the lengths

of the four independent coils, from $L_{\max} = 18\text{ m}$ to $L_{\max} = 24\text{ m}$ to obtain several coil sets of varying complexity and performance. These lengths are relative to a device configuration scaled to have an average major radius of 1 m. We refer to a specific coilset using the notation ‘Target[L_{\max}]’, e.g., QA[24] refers to the coilset that targets the QA configuration of [16] and for which the four independent modular coils have a combined length of maximally 24 m.

As demonstrated in [26], solving this optimization problem yields coils that achieve very precise quasi-symmetry and low particle losses. However, for the design of a practical experiment or reactor it is crucial to understand the impact that coil errors have on the magnetic field and on its ability to confine particles. In order to answer this question, we require a model for coil errors that arise during manufacturing and placement. Following the approach of [27, Section 2], we denote the coil errors by $\Xi^{(i)}$ and consider the perturbed coils with parameterization $\Gamma^{(i)} + \Xi^{(i)}$. We model coil errors as the sum of a systematic error, which satisfies two-field-period symmetry and stellarator symmetry, and statistical error, which is independent for each coil, that is $\Xi^{(i)} = \Xi_{\text{sys}}^{(i)} + \Xi_{\text{stat}}^{(i)}$. The two different sources of error are modelled using a Gaussian process with the same pointwise standard deviation σ . To give intuition for the resulting magnitude of the error, we consider the normal component of the perturbation. Denoting by \mathbf{n} and \mathbf{b} a normal and binormal vector along the curve, we note that at each point, $P(\theta) = \|(\Xi^{(i)}(\theta) \cdot \mathbf{n}(\theta))\mathbf{n}(\theta) + (\Xi^{(i)}(\theta) \cdot \mathbf{b}(\theta))\mathbf{b}(\theta)\|$ is Rayleigh($\sqrt{2}\sigma$) distributed, and hence the expected pointwise normal error magnitude is $\mathbf{E}(P(\theta)) = \sqrt{\pi}\sigma$. The distribution of the maximum error along the coil is more complicated, but in our experiments it is usually a small multiple of the expected normal error magnitude.

The length scales for the stochastic processes are chosen to be $\pi/2$ for the systematic error, and π for the statistical error. An example for a coil perturbation drawn from this model can be seen in Figure 2; in magnitude and structure this error is similar to that measured for the W7-X experiment [2, Figure 3].

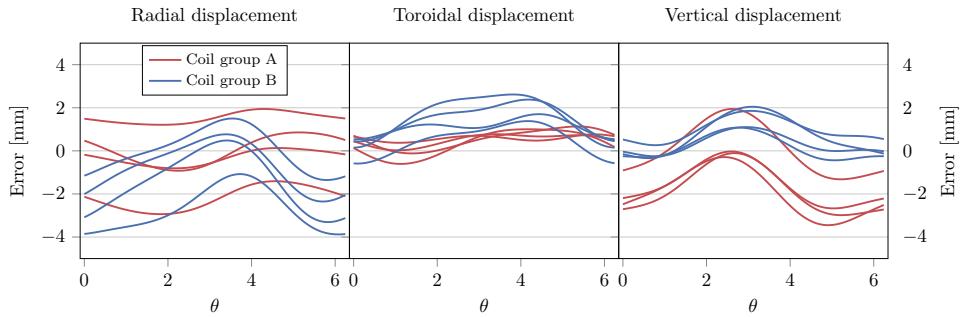


Figure 2. Samples drawn from the error model for two coil groups (A and B). Shown in the same colour are samples from our error model for coils related through stellarator and two-field-period symmetry. We see that part of the error is the same across the four instances of the modular coils, i.e., within group A and group B. The reader may compare these figures with the measurements for W7-X in [2, Figure 3].

We collect the information about the Gaussian processes underlying the deformations in ζ and denote the magnetic field generated by a set of perturbed

coils by

$$\mathbf{B}(\mathbf{c}, \zeta) = \mathbf{B}(\{\Gamma^{(i)}(\mathbf{c}) + \Xi^{(i)}(\zeta)\}_{i=1}^{16}, \{I_i(\mathbf{c})\}_{i=1}^{16}). \quad (2)$$

The stochastic optimization problem that we solve is then given by

$$\underset{\mathbf{c}}{\text{minimize}} \quad \mathbf{E}_{\zeta} \left[f(\mathbf{B}(\mathbf{c}, \zeta)) \right], \quad (3a)$$

$$\text{subject to} \quad \text{Max}(\kappa_i) \leq \kappa_{\max} \text{ for } i = 1, \dots, 4, \quad (3b)$$

$$\text{Mean}(\kappa_i^2) \leq \kappa_{\text{msc}} \text{ for } i = 1, \dots, 4, \quad (3c)$$

$$\sum_{i=1}^4 L_i \leq L_{\max}, \quad (3d)$$

$$\min_{\theta, \theta'} \|\Gamma_i(\theta) - \Gamma_j(\theta')\| \geq d_{\min} \text{ for all } 1 \leq i < j \leq 16, \quad (3e)$$

$$\|\Gamma'_i\| = \frac{L_i}{2\pi} \text{ for } i = 1, \dots, 4. \quad (3f)$$

Here (3b) limits the maximum curvature along each coil, (3c) limits the mean-squared curvature along each coil, (3d) limits the combined length of the independent modular coils, and (3e) enforces that the coils are separated from each other. Finally, the constraint in (3f) ensures a constant incremental arclength in the parametrization of the coils. We include this constraint for two reasons:

- (i) To ensure differentiability of line integrals: the mean squared curvature for example is given by $\frac{1}{L_i} \int_{[0, 2\pi)} \kappa_i(\theta) \|\Gamma'_i(\theta)\| d\theta$. The incremental arclength term in that integral is not differentiable with respect to coil coefficients if $\|\Gamma'_i(\theta)\| = 0$ anywhere.
- (ii) To ensure constant length scale of the Gaussian process error model: if the incremental arclength is allowed to vary along the coil, then this might distort the characteristic length scale of the Gaussian process. In experiments without this constraint, we observe that the optimization algorithm exploits this freedom. This results in a lower objective value, deceptively implying that stochastic optimization has larger benefit than it actually does. In reality, some of the improvement in the objective is due to varying the arclength around the curve. This causes quadrature points along the coil to bunch together or to separate from one another, thereby nonphysically smoothing the error in some regions, or making it highly oscillatory in others. This is not desirable as the representation of the curve should not impact the nature of the errors.

To approximate the expectation over the coil errors, we draw samples ζ_k of the error distribution, evaluate the flux objective for each perturbed set of coils and compute the Monte Carlo sample average, that is

$$\mathbf{E}_{\zeta} [f(\mathbf{B}(\mathbf{c}, \zeta))] \approx \frac{1}{N} \sum_{k=1}^N f(\mathbf{B}(\mathbf{c}, \zeta_k)). \quad (4)$$

The deterministic optimization problem can be recovered by simply setting $\zeta = 0$.

We solve this optimization problem using the SIMSOPT software [15]. SIMSOPT is an open-source library for stellarator optimization that is mostly written in Python, with performance-critical parts (such as the Biot-Savart law) implemented in C++. The constraints are implemented as penalty functions, and we increase each penalty

term until the constraint violation is less than 0.1%, as also done in [26]. Derivatives of the flux objective and the regularization terms are computed by SIMSOPT using a mix of manual implementation and automatic differentiation based on JAX [6]. SIMSOPT uses the L-BFGS [22] quasi-Newton scheme to solve the optimization problem efficiently. L-BFGS uses a history of gradients evaluated at different parameters to approximate second derivatives, which we found to be important for convergence as it helps to overcome the ill-posedness of the optimization problem. We run the L-BFGS iteration for 14 000 iterations and then perform an additional five steps of Newton's method, using a finite difference approximation for the Hessian, in order to ensure convergence. For the shorter coils the problem is better conditioned and L-BFGS on its own is able to reduce the norm of the gradient of the objective by 10 orders of magnitude or more. For the longer coils, L-BFGS is only able to reduce the gradient by about seven orders of magnitude and we gain an extra order of magnitude reduction using the Newton iterations. We note that poor conditioning of coil design problems as the coil length increases was also observed in [7, Fig. 4].

To mitigate issues of local minima, we ran the deterministic optimization for 8 different initial guesses and picked the best performing configuration, as measured by the flux objective. The different initial guesses are obtained by performing random perturbations of regularly spaced circular coils. For the stochastic optimization, we run 16 optimizations: 8 that use the minimizers of the deterministic problem as initial guesses, and 8 that use perturbations of circular coils as initial guess. We use 4096 samples to approximate the mean in (4). Each stochastic optimization problem is run on 128 cores across eight nodes consisting of two Intel Xeon Platinum 8268 CPUs and takes approximately eight hours to solve.

2.2. Numerical results

First, we evaluate the impact of stochastic optimization directly on the flux objective. Then, we additionally study how designs obtained with the deterministic and the stochastic approaches perform with respect to physical quantities of interest such as alpha particle confinement.

We solve the stochastic problems for perturbations with $\sigma = 1\text{ mm}$, as well as the deterministic problem. For the obtained coil designs, we then draw 128 new perturbations of magnitude ranging from $\sigma = 0\text{ mm}$ to $\sigma = 1\text{ mm}$ and compute the sample mean of the flux objective. Here, and in the remainder of this manuscript, we use the notation $\langle \cdot \rangle$ to denote that we computed the sample average across coil perturbations for the quantity of interest (such as quadratic flux, violation of quasi-symmetry, or particle confinement). We recall that the device has a major radius of approximately $R = 1\text{ m}$ and that $\sigma = 1\text{ mm}$ implies an expected pointwise error of $\sqrt{\pi}\sigma \approx 0.002R$. For devices with different major radius R , the coil error magnitude would be scaled linearly to maintain the same relative quadratic flux 1.

As expected, for very small perturbations the deterministic minima outperform the stochastic minima. If this was not the case, this would indicate that the deterministic algorithm was stuck in a poor local minimum. We only observe an advantage for stochastic optimization over deterministic optimization once the perturbation magnitude is close to the one used for the stochastic optimization. However, the difference is not large: for the QA[24] and QA+Well[24] configurations the mean is reduced by $\sim 13\%$ and $\sim 10\%$ respectively, and the improvement is smaller still for the shorter coils.

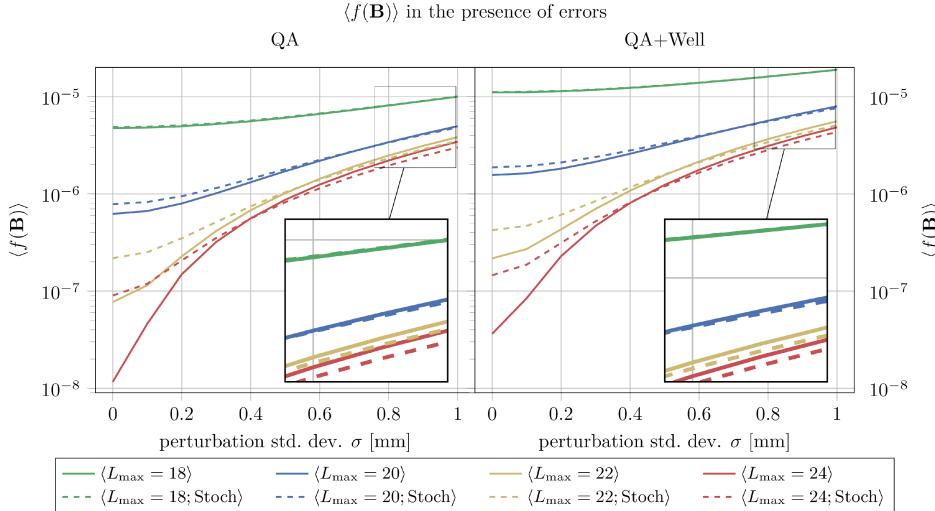


Figure 3. Impact of coil errors of different magnitude σ (shown on the x -axis) on the average flux $\langle f(\mathbf{B}) \rangle \equiv \mathbf{E}_\zeta[f(\mathbf{B}(\mathbf{c}, \zeta))]$ for deterministic and stochastic coil designs. The stochastic coil optimization used an error model with $\sigma = 1$ mm. The magnified detail highlights the improved performance of coils from stochastic optimization in the presence of errors, though this improvement is small.

We recall that the motivation behind the flux objective in (1) is that if $\mathbf{B} \cdot \mathbf{n} = 0$ everywhere on the surface, then the field generated by the coils matches the target field everywhere inside the volume surrounded by the surface. In particular, it means that the field induced by the coils has the same confinement and quasi-symmetry properties. However, we are not guaranteed that there is a strict correlation between the flux objective and the physics properties for non-zero values of $f(\mathbf{B})$. This means that we need to evaluate the quasi-symmetry and the confinement performance of the perturbed coils to study to what extent an improvement of $\langle f(\mathbf{B}) \rangle$ translates to improvements in the physical quantities of interest.

To study this, we compute a measure of quasi-symmetry and particle confinement for 128 different coil perturbations and compare the performance to that achieved when building and placing the coils obtained from deterministic optimization exactly. To quantify quasi-symmetry, we recall that a magnetic field is quasi-axisymmetric if the field strength is independent of the angle θ when expressed in terms of Boozer angles φ and θ [12]. Thus we can quantify the deviation from quasi-symmetry as follows: for each perturbed coilset, with currents scaled so the average field strength is 1 T, we compute surfaces in Boozer angles for the magnetic field using [7], and perform a Fourier transform of the field strength. That is, we express $|\mathbf{B}(s, \theta, \varphi)| = \sum_{m,n} B_{m,n}(s) \cos(m\theta - n\varphi)$. For each flux value s , we report the largest symmetry-breaking component, i.e. $\max_{n \neq 0, m} |B_{m,n}(s)|$, and then compute the sample average over all coil perturbations. The result is shown in Figure 4. The coil errors result in a significant departure from quasi-axisymmetry. For most, but not all of the configurations, the perturbed stochastic minimizers have slightly better quasi-symmetry than the perturbed deterministic configurations. We emphasize, however, that this improvement pales in comparison to the significant departure from quasi-symmetry due to coil errors.

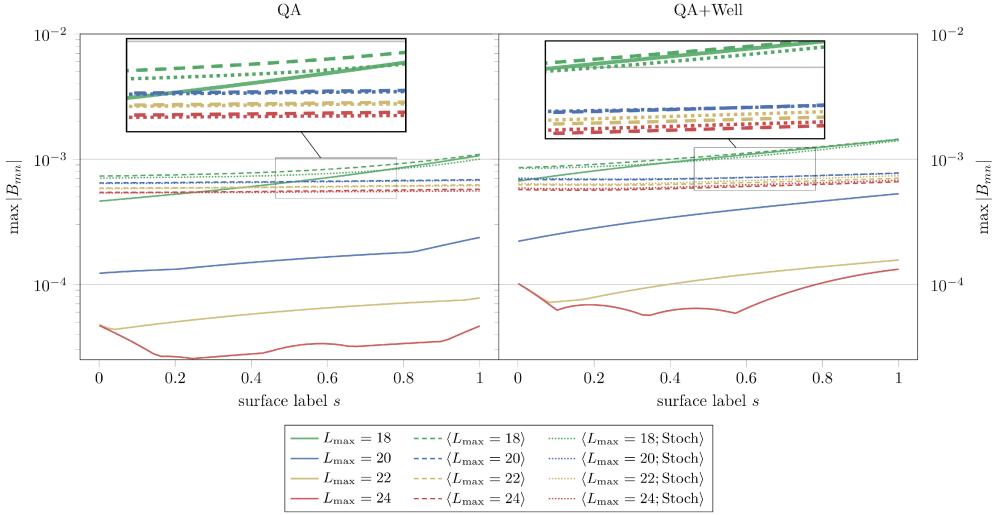


Figure 4. Sample average of the largest quasi-symmetry breaking modes for $\sigma = 1\text{ mm}$. Solid lines: coils from deterministic optimization built exactly. Dashed lines: sample average for coils from deterministic optimization with perturbations. Dotted lines: sample average for coils from stochastic optimization with perturbations.

Finally, we evaluate the ability of the fields to confine highly energetic particles. To this end, we scale the device to approximately match the ARIES-CS [20] reactor: we choose a minor radius of 1.7 m and an average field strength on the outer surface of 5.86 T. We spawn alpha particles at the surface with half radius and initialize them isotropically in velocity space with 3.5 MeV energy, and numerically compute their collisionless guiding-center trajectories with the SIMSOPT code [15]. Again we consider the case in which the coils obtained from deterministic optimization were built and placed perfectly as the reference case, which we compare with a sample average over 128 perturbed coil configurations for both the deterministic and stochastic optimization case. The losses are shown in Figure 5. We make two main observations: firstly, compared to the case when coils are built exactly (solid lines), the perturbed coils have significantly higher energetic particle losses. And secondly, the coils found by stochastic optimization do not systematically outperform those from deterministic optimization. For the shorter coils, stochastic optimization seems to result in slightly better confinement, but the opposite is true for the longer coils. The conclusion here is that when deterministic coil design is done in a way that avoids being stuck in poor local minima, stochastic optimization does not yield a substantial benefit for the classical two-stage coil design problem, at least for the vacuum fields we have considered in this work.

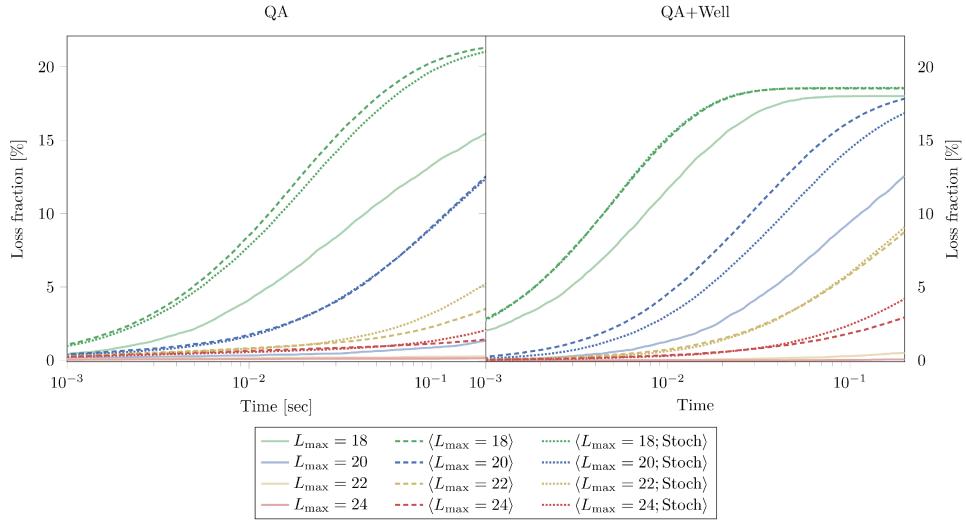


Figure 5. Confinement of alpha particles spawned on the $s = 0.25$ surface. Solid lines: coils from deterministic optimization built exactly. Dashed lines: sample average for coils from deterministic optimization with perturbations. Dotted lines: sample average for coils from stochastic optimization with perturbations.

2.3. Discussion of the impact of stochastic optimization

Stochastic optimization has received significant attention over the last few years in the stellarator design community. We focus on the common case of risk-neutral stochastic optimization, i.e., we consider expectations over the random parameter distribution. In order to put our results into context with other work, we start by recalling two main ways in which stochastic optimization impacts objective functions and their minimizers. As an illustrative example, we assume that the stochastic objective is of the form $\mathbf{E}[f(\mathbf{x} + \boldsymbol{\xi})]$ for a Gaussian random vector $\boldsymbol{\xi}$. In this case, stochastic optimization can be interpreted as a convolution with a Gaussian kernel. In particular, stochastic objectives typically are smooth versions of their deterministic counterparts. This has two main implications, which are demonstrated in Figure 6. Firstly, this smoothing may reduce the number of local minima and maxima. Secondly, deterministic minima that are very close to much larger function values may not be minima of the stochastic objective as the stochastic objective takes into account a neighborhood around each point. This is why minimizers computed from stochastic formulations tend to perform robustly with respect to random perturbations. Since we consider several different initial guesses for the optimization and pick the best performing local minimum, any difference between deterministic and stochastic optimization in our results are likely due to the second effect above.

Stochastic optimization was previously applied to the problem of coil design in [19], where a stochastic version of the ONSET code is used to find coils for the W7-X experiment and it is observed that the objective value is reduced by up to 21% when comparing stochastic to deterministic minimization. The main reason for this significant improvement appears to be due to the deterministic minimizer getting stuck in a poor local minimum. We can see that this is the case, since even in the absence of

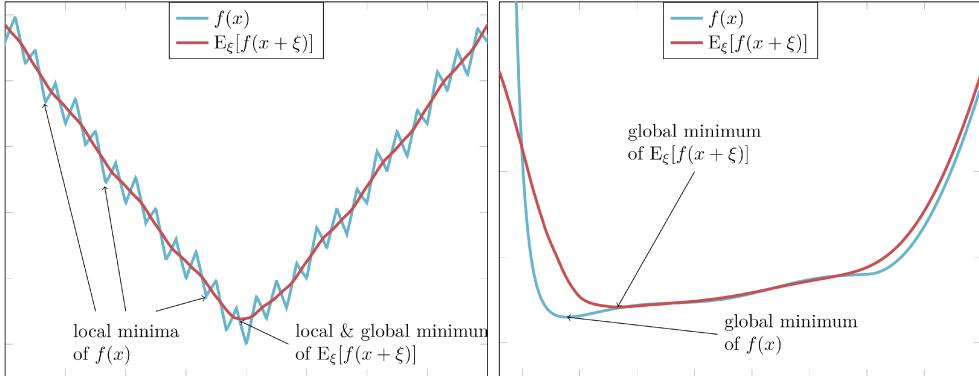


Figure 6. Illustration of effects of stochastic optimization. Shown on the left is the fact that stochastic optimization tends to smooth objective functions and on the right that it tends to find minimizers that are more stable, i.e., they lie in wider basins of the original function.

errors, the minima obtained from stochastic optimization outperform those obtained from deterministic optimization by up to 19%—implying that the latter can not be the global minimum (see Figure 1 in [19] and compare the values given for $f(x_0)$). Hence, it seems to be the first of the two effects in Figure 6 that has a significant impact.

We postulate that the reason why ONSET struggles with local minima is a combination of the choice of objective and optimization algorithm: ONSET contains several non-smooth contributions to the objective (such as the maximal field error), and uses a derivative-free optimization algorithm. In contrast, the FOCUS-like formulation that we use here only involves smooth functions and we use a derivative-based algorithm, which appears to lead to fewer problems with local minima. In fact, we note that the different local minima that we find as a result of using different initialization all perform rather similarly well and, while different, are visually similar too.

More recently, in [27] stochastic optimization was applied to single-stage coil design, that directly optimized coils for quasi-symmetry near the magnetic axis, as presented in [8]. When performing deterministic optimization for the single-stage quasi-symmetry objective in [8], we are able to achieve objective values of $< 10^{-10}$. However, even small perturbations result in objective values that are several orders of magnitude larger. This suggests that we are in the second situation of Figure 6, and that the improvement is not due to the avoidance of local minima, but due to stochastic minimization avoiding unstable locations in parameter space.

3. A posteriori coil error correction

Much of today’s understanding of coil manufacturing errors and their impact can be credited to the Wendelstein 7-X experiment. During the construction of W7-X, manufacturing errors of the coils were measured after their production and the position of the coils was adjusted one by one [3, 5] to minimize error-fields to first order. Motivated by this idea, we now systematically study the ability of a-posteriori

coil adjustments to reduce the impact of coil errors. We will use the error model from the previous section to simulate coil errors of various magnitudes and then evaluate how well these errors can be mitigated.

While, once built, the shape of the coil cannot be changed, the coil can still be shifted, rotated, and its current can be adjusted. There are therefore seven degrees of freedom per coil that could be optimized: three degrees of freedom corresponding to translations in each direction, three degrees of freedom corresponding to rotations around each axis, and one degree of freedom for the current in the coil. Here, we consider each of the 16 coils independently, as manufacturing errors might not satisfy stellarator symmetry or field period symmetry and hence each coil might require a distinct correction.

Denoting, for the i th coil, the translation by $\mathbf{s}^{(i)} \in \mathbb{R}^3$ and the rotation angles by $\boldsymbol{\alpha}^{(i)} = (\alpha_1^{(i)}, \alpha_2^{(i)}, \alpha_3^{(i)}) \in \mathbb{R}^3$, we consider the rigid transformation

$$\mathbf{O}_{\mathbf{s}^{(i)}, \boldsymbol{\alpha}^{(i)}}(\mathbf{x}) = \mathbf{R}_{\text{yaw}}(\alpha_1^{(i)}) \mathbf{R}_{\text{pitch}}(\alpha_2^{(i)}) \mathbf{R}_{\text{roll}}(\alpha_3^{(i)}) \mathbf{x} + \mathbf{s}^{(i)} \quad (5)$$

with the rotation matrices

$$\begin{aligned} \mathbf{R}_{\text{yaw}}(\alpha) &= \begin{bmatrix} \cos(\alpha) & -\sin(\alpha) & 0 \\ \sin(\alpha) & \cos(\alpha) & 0 \\ 0 & 0 & 1 \end{bmatrix}, \\ \mathbf{R}_{\text{pitch}}(\alpha) &= \begin{bmatrix} \cos(\alpha) & 0 & \sin(\alpha) \\ 0 & 1 & 0 \\ -\sin(\alpha) & 0 & \cos(\alpha) \end{bmatrix}, \\ \mathbf{R}_{\text{roll}}(\alpha) &= \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos(\alpha) & -\sin(\alpha) \\ 0 & \sin(\alpha) & \cos(\alpha) \end{bmatrix}. \end{aligned} \quad (6)$$

We denote the magnetic field induced by the corrected coils by

$$\mathbf{B}(\mathbf{c}, \boldsymbol{\zeta}, \{(\mathbf{s}^{(i)}, \boldsymbol{\alpha}^{(i)}, \delta I^{(i)})\}_{i=1}^{16}) = \mathbf{B}(\{\mathbf{O}_{\mathbf{s}^{(i)}, \boldsymbol{\alpha}^{(i)}}(\boldsymbol{\Gamma}^{(i)}(\mathbf{c}) + \boldsymbol{\Xi}^{(i)}(\boldsymbol{\zeta}))\}_{i=1}^{16}, \{I^{(i)}(\mathbf{c}) + \delta I^{(i)}\}_{i=1}^{16}) \quad (7)$$

and then for a given design $\hat{\mathbf{c}}$ (obtained from deterministic optimization in this case) and measured manufacturing errors $\hat{\boldsymbol{\zeta}}$, we solve

$$\begin{aligned} &\underset{\{(\mathbf{s}^{(i)}, \boldsymbol{\alpha}^{(i)}, \delta I^{(i)})\}_{i=1}^{16}}{\text{minimize}} \quad f(\mathbf{B}(\hat{\mathbf{c}}, \hat{\boldsymbol{\zeta}}, \{(\mathbf{s}^{(i)}, \boldsymbol{\alpha}^{(i)}, \delta I^{(i)})\}_{i=1}^{16})), \\ &\text{subject to} \quad \begin{cases} \text{the curves } \{\mathbf{O}_{\mathbf{s}^{(i)}, \boldsymbol{\alpha}^{(i)}}(\boldsymbol{\Gamma}^{(i)}(\mathbf{c}) + \boldsymbol{\Xi}^{(i)}(\boldsymbol{\zeta}))\}_{i=1}^{16} \\ \text{satisfying the distance constraints.} \end{cases} \end{aligned} \quad (8)$$

Since coil length and curvature are not affected by rigid transformations, all other geometric constraints considered in the optimization (3) or its deterministic counterpart are satisfied automatically. We note that our approach differs from the one employed for W7-X in a few aspects. Firstly, all coils are adjusted at once, rather than sequentially. Secondly, the objective that we minimize is the same as the one used during coil design, rather than a separate error field objective, and lastly, in addition to considering coil shifts and rotations, we also consider current adjustments as a degree of freedom that can be used to adjust the magnetic field. To better understand the impact of the current in addition to the placement correction, we consider three variants of this a-posteriori optimization, which we refer to as the correction levels (CL):

- (CL1) Only optimize for the position and orientation of the coils. The resulting optimization problem has $16 \cdot 6 = 96$ degrees of freedom.
- (CL2) CL1 plus a current correction per modular coil group. This is motivated by the fact that usually each set of four nominally identical modular coils is connected in series to the same power supply, so adjusting the current on a per-coil basis is more complicated. The resulting optimization problem has $16 \cdot 6 + 4 = 100$ degrees of freedom.
- (CL3) CL1 plus a current correction for each individual coil. The resulting optimization problem has $16 \cdot 7 = 112$ degrees of freedom. This corresponds to the situation in (8).

In Figure 7, we show an example of the impact of coil errors and how well they can be corrected. While in the figure we only show four coils, all 16 coils may have different errors, and each coil position is corrected by solving the optimization problem (8) over rigid body motions and currents for each coil. As can be seen, the normal flux increases significantly due to errors, but that increase can be almost entirely eliminated by the coil position, orientation, and current adjustments.

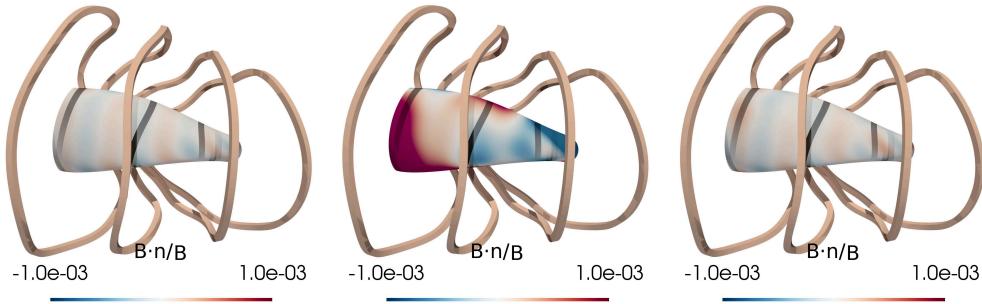


Figure 7. Left: QA[22] coils. Middle: QA[22] coils with random draw from the error model ($\sigma = 2 \text{ mm}$). Right: QA[22] coils with error and CL3 correction. Note that the differences between the coils are too small to be visible. Coils are modelled as line filaments and the finite cross section is for visual purposes only. The color on the surface indicates $\mathbf{B} \cdot \mathbf{n} / |\mathbf{B}|$. We observe that the a-posteriori corrections are able to almost entirely undo the impact of the coil errors.

Remark. Note that when building a stellarator experiment, one may not be able to wait until all the coils have been manufactured to start placing them. To address such situations, one can formulate a sequence of similar optimization problems that adjust coils or groups of coils as they arrive from manufacturing. The all-at-once optimization problem considered here can be viewed as a best case scenario in which measurements of all coils are available.

To systematically evaluate to what degree accurately measured manufacturing errors can be corrected, we again draw 128 random perturbations to be added to the deterministic minimizer. Then, for each set of perturbations, we solve the optimization problem (8) for each of the correction levels in order to minimize the flux objective. We repeat this procedure for multiple perturbations and error magnitudes and plot the resulting sample averages in Figure 8. We find that coil manufacturing errors can consistently be corrected to a large degree by adjusting coil placement and currents.

Even when keeping the currents fixed and only adjusting for the position of the coils, the objective value only increases slightly, i.e., most of the manufacturing error can be corrected by a slight correction of the coil position. In particular for the shorter coils, (CL1), (CL2), and (CL3) yield near identical results. For the longest coils and largest errors considered here, adding group-wise current corrections (CL2) improves the flux objective by $\sim 18\%$, and adding individual current corrections (CL3) reduces the error by another $\sim 4\%$.

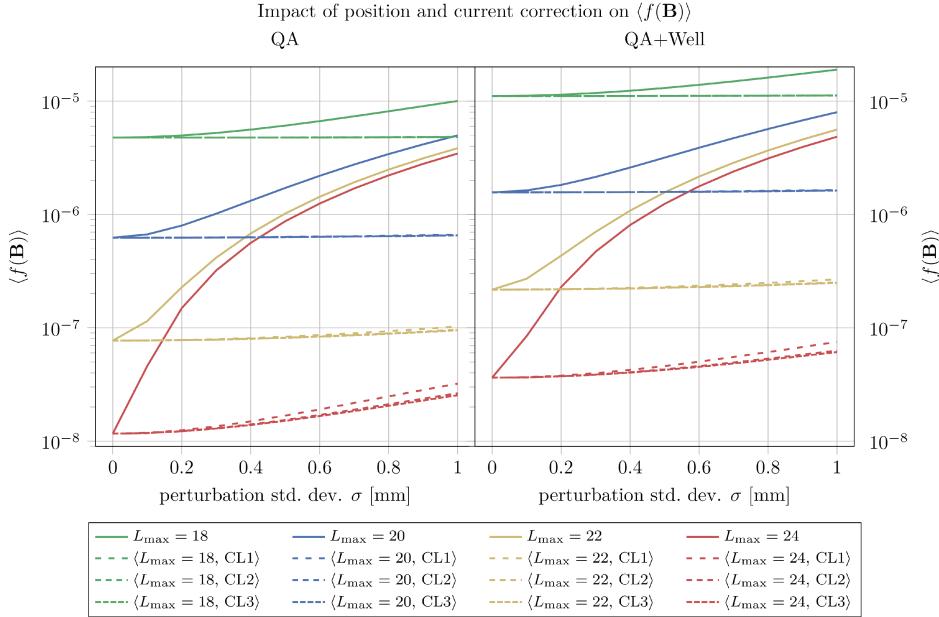


Figure 8. Impact of coil errors with different σ (x-axis) on $\langle f(\mathbf{B}) \rangle$ for the deterministic minimizers both without correction (solid) and with the three different a posterior correction levels (dashed). It can be seen that the corrections substantially mitigate the degradation of the flux objective.

To study how well improvement in the flux objective translates to better physics performance, we compute the quasi-symmetry violation and the energetic particle confinement properties for 128 perturbed stellarators and for each correction level. For conciseness, we focus on the case $L_{\max} = 24$ and $\sigma = 1$ mm for these computations. In Figure 9 we can see that the position correction (CL1) yields a significant improvement of quasi-symmetry. Adding the current correction in CL2 and CL3 then yields incremental additional improvement.

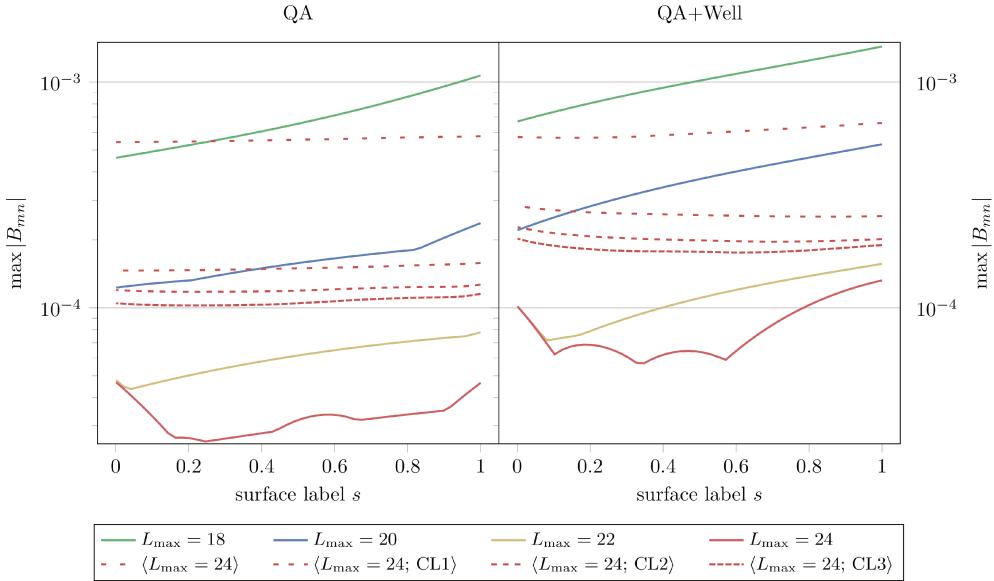


Figure 9. Sample average for errors with $\sigma = 1$ mm of the largest quasi-symmetry breaking modes as a function of surface label s (x -axis). Shown are the results for the longest coilset without correction and with all three correction levels (dashed lines). As reference, we also show the results without errors (solid lines) for different coil lengths L_{\max} .

Finally, using the same setup as described at the end of Section 2.2, we consider alpha particle confinement in Figure 10. Without corrections, we lose $\sim 1.42\%$ and $\sim 2.92\%$ of all particles respectively after 200 ms. We observe that particle losses are reduced significantly when adjusting the position of the coils, almost to the magnitude obtained by coils without manufacturing errors. Furthermore, minor improvement can be obtained by also adjusting their currents. For the highest correction level, we lose only $\sim 0.26\%$ for the QA, and $\sim 0.17\%$ of the QA+Well configuration, a fairly modest increase compared to $\sim 0.14\%$ and $\sim 0.06\%$ respectively for the perfectly built and placed coils.

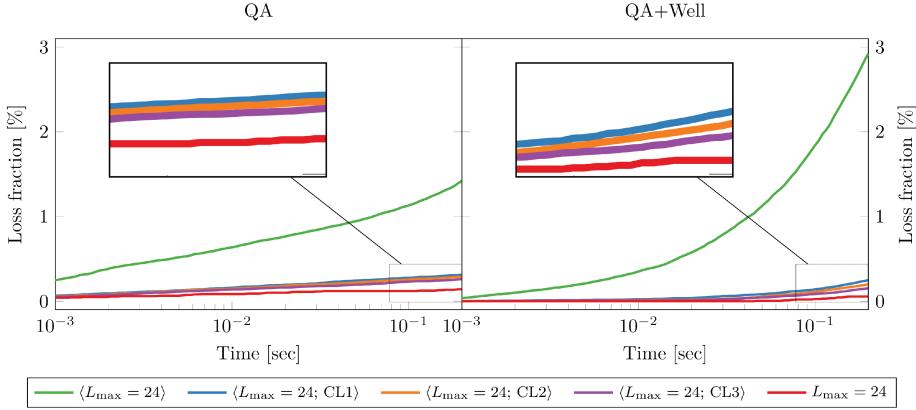


Figure 10. Particle confinement of alpha particles spawned on the $s = 0.25$ surface for the exactly built coils (red), for 128 perturbed coils without correction (green), and with all three correction levels (blue, orange, purple). As the correction level increases, particle losses decrease.

4. Discussion and Conclusion

In this work we have investigated the impact of coil manufacturing errors on stellarator coils generating magnetic fields that are quasi-symmetric to a very high accuracy. We have shown that fairly small errors have considerable impact on the confinement properties of the induced magnetic fields, and that stochastic optimization does not enable substantial mitigation of coil errors. On the other hand, if we assume that coil errors can be measured accurately, we showed that particle confinement can be improved significantly by solving a separate optimization problem for an optimal adjustment of the coils. This adjustment includes shifts and rotations of the coils, and optionally can include changes of the coil currents. While we used our error model to assess this approach, we highlight that no such model is required when applying this approach during construction: the initial coil design can be performed without consideration of errors, and the a-posteriori optimization corrects for specific manufacturing errors after they were measured.

While, in practice, errors cannot be perfectly measured and coils cannot be placed perfectly, our study highlights the efficacy of a-posteriori coil adjustments as a tool for dealing with coil manufacturing errors.

Code availability

The code used for the simulations and analysis in this manuscript is publicly available at

<https://github.com/florianwechsung/Stage-II-Optimization-With-Perturbations>
and has been archived at [25].

Acknowledgments

The authors would like to thank the SIMSOPT development team and Thomas Sunn Pedersen for his insight on the construction and coil errors of W7X. In addition,

this work was supported in part through the NYU IT High Performance Computing resources, services, and staff expertise. This work was supported by a grant from the Simons Foundation (560651). AG is partially supported by an NSERC (Natural Sciences and Engineering Research Council of Canada) postdoctoral fellowship. In addition, AC and FW are supported by the United States National Science Foundation under grant No. PHY-1820852, and AC is supported by the United States Department of Energy, Office of Fusion Energy Sciences, under grant No. DE-FG02-86ER53223.

References

- [1] T. Andreeva et al. “Analysis of the Magnetic Field Perturbations during the Assembly of Wendelstein 7-X”. In: *Fusion Science and Technology* 46.2 (Sept. 2004), pp. 388–394. DOI: 10.13182/FST04-A578.
- [2] T. Andreeva et al. “Final Assessment of Wendelstein 7-X Magnetic Field Perturbations Caused by Construction Asymmetries”. In: *25th IAEA Fusion Energy Conference* (2014), p. 7.
- [3] T. Andreeva et al. “Influence of construction errors on Wendelstein 7-X magnetic configurations”. In: *Fusion Engineering and Design* 84.2-6 (June 2009), pp. 408–412. ISSN: 09203796. DOI: 10.1016/j.fusengdes.2008.12.073. (Visited on 03/15/2022).
- [4] A. Bader et al. “Advancing the physics basis for quasi-helically symmetric stellarators”. In: *Journal of Plasma Physics* 86.5 (Oct. 2020), p. 905860506. ISSN: 0022-3778, 1469-7807. DOI: 10.1017/S0022377820000963.
- [5] H.-S. Bosch et al. “Technical challenges in the construction of the steady-state stellarator Wendelstein 7-X”. In: *Nuclear Fusion* 53.12 (Dec. 2013), p. 126001. ISSN: 0029-5515, 1741-4326. DOI: 10.1088/0029-5515/53/12/126001.
- [6] James Bradbury et al. *JAX: composable transformations of Python+NumPy programs*. Version 0.2.5. 2018. URL: <http://github.com/google/jax>.
- [7] Andrew Giuliani et al. “Direct computation of magnetic surfaces in Boozer coordinates and coil optimization for quasi-symmetry”. In: *arXiv preprint arXiv:2203.03753* (2022).
- [8] Andrew Giuliani et al. “Single-stage gradient-based stellarator coil design: Optimization for near-axis quasi-symmetry”. In: *Journal of Computational Physics* 459 (2022), p. 111147.
- [9] Silke Glas et al. “Global stochastic optimization of stellarator coil configurations”. In: *Journal of Plasma Physics* 88.2 (2022).
- [10] K. C. Hammond et al. “Experimental and numerical study of error fields in the CNT stellarator”. In: *Plasma Physics and Controlled Fusion* 58.7 (July 2016), p. 074002. DOI: 10.1088/0741-3335/58/7/074002.
- [11] C.C. Hegna et al. “Improving the stellarator through advances in plasma theory”. In: *Nuclear Fusion* 62.4 (Apr. 2022), p. 042012. DOI: 10.1088/1741-4326/ac29d0.
- [12] Per Helander. “Theory of plasma confinement in non-axisymmetric magnetic fields”. In: *Reports on Progress in Physics* 77.8 (July 2014), p. 087001. DOI: 10.1088/0034-4885/77/8/087001.

- [13] JCT ITER. “Technical basis for the ITER final design report, cost review and safety analysis (FDR)”. In: *ITER/EDA Documentation Series, IAEA, Vienna* (1999).
- [14] T. Klinger et al. “Towards assembly completion and preparation of experimental campaigns of Wendelstein 7-X in the perspective of a path to a stellarator fusion power plant”. In: *Fusion Engineering and Design* 88.6 (2013). Proceedings of the 27th Symposium On Fusion Technology (SOFT-27); Liège, Belgium, September 24-28, 2012, pp. 461–465. ISSN: 0920-3796. DOI: <https://doi.org/10.1016/j.fusengdes.2013.02.153>.
- [15] M. Landreman et al. “SIMSOPT: A flexible framework for stellarator optimization”. In: *Journal of Open Source Software* 6 (2021), p. 3525.
- [16] Matt Landreman and Elizabeth Paul. “Magnetic fields with precise quasisymmetry for plasma confinement”. In: *Physical Review Letters* 128.3 (2022), p. 035001.
- [17] Jim-Felix Lobsien et al. “Improved performance of stellarator coil design optimization”. In: *Journal of Plasma Physics* 86.2 (Apr. 2020). DOI: 10.1017/S0022377820000227.
- [18] Jim-Felix Lobsien et al. “Physics analysis of results of stochastic and classic stellarator coil optimization”. In: *Nuclear Fusion* 60.4 (Apr. 2020). DOI: 10.1088/1741-4326/ab7211.
- [19] Jim-Felix Lobsien et al. “Stellarator coil optimization towards higher engineering tolerances”. In: *Nuclear Fusion* 58.10 (Oct. 2018). DOI: 10.1088/1741-4326/aad431.
- [20] F. Najmabadi et al. “The ARIES-CS Compact Stellarator Fusion Power Plant”. In: *Fusion Science and Technology* 54.3 (Oct. 2008), pp. 655–672. ISSN: 1536-1055, 1943-7641. DOI: 10.13182/FST54-655.
- [21] G. H. Neilson et al. “Lessons Learned in Risk Management on NCSX”. In: *IEEE Transactions on Plasma Science* 38.3 (2010), pp. 320–327.
- [22] Jorge Nocedal. “Updating Quasi-Newton Matrices with Limited Storage”. In: *Mathematics of Computation* 35.151 (July 1980), p. 773. ISSN: 00255718. DOI: 10.2307/2006193. (Visited on 04/25/2019).
- [23] Akihiro Shimizu et al. “Consideration of the Influence of Coil Misalignment on the Chinese First Quasi-Axisymmetric Stellarator Magnetic Configuration”. In: *Plasma and Fusion Research* 14.0 (Sept. 2019), pp. 3403151–3403151. ISSN: 1880-6821. DOI: 10.1585/pfr.14.3403151. (Visited on 03/18/2022).
- [24] R. L. Strykowsky et al. “Engineering cost schedule lessons learned on NCSX”. In: *2009 23rd IEEE/NPSS Symposium on Fusion Engineering*. 2009, pp. 1–4.
- [25] Florian Wechsung. *Stage-II-Optimization-With-Perturbations: v1.0*. Mar. 2022. DOI: 10.5281/zenodo.6369168. URL: <http://dx.doi.org/10.5281/zenodo.6369168>.
- [26] Florian Wechsung et al. “Precise stellarator quasi-symmetry can be achieved with electromagnetic coils”. In: *Proceedings of the National Academy of Sciences* 119.13 (2022), e2202084119.
- [27] Florian Wechsung et al. “Single-stage gradient-based stellarator coil design: stochastic optimization”. In: *Nuclear Fusion* 62.7 (2022), p. 076034.

- [28] C. Windsor et al. “Alpha-particle physics in tokamaks”. In: *Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences* 357.1752 (1999), pp. 493–513. DOI: 10.1098/rsta.1999.0338. eprint: <https://royalsocietypublishing.org/doi/pdf/10.1098/rsta.1999.0338>. URL: <https://royalsocietypublishing.org/doi/abs/10.1098/rsta.1999.0338>.
- [29] Caoxiang Zhu et al. “New method to design stellarator coils without the winding surface”. In: *Nuclear Fusion* 58.1 (Jan. 2018), p. 016008. ISSN: 0029-5515, 1741-4326. DOI: 10.1088/1741-4326/aa8e0a.