Stellarator coil optimization towards higher engineering tolerances

Jim-Felix Lobsien

Max-Planck Institute for Plasma Physics, Greifswald, Germany

E-mail: jim.lobsien@ipp.mpg.de

Michael Drevlak

Max-Planck Institute for Plasma Physics, Greifswald, Germany

E-mail: michael.drevlak@ipp.mpg.de

Thomas Sunn Pedersen

Max-Planck Institute for Plasma Physics, Greifswald, Germany

E-mail: thomas.sunn.pedersen@ipp.mpg.de

W7-X Team

Abstract. Recently designed optimized stellarator experiments have suffered from very tight construction tolerances, but some level of deviation of the coil system is unavoidable during fabrication of the coils and assembly of the coil system. In this paper, we present a new approach that incorporates reduced sensitivity to construction tolerances of the coil system into the optimization sequence. The approach was tested within the framework of the existing coil optimization scheme for Wendelstein 7-X. The results are compared with those of a coil set obtained by the original optimization. The result is a more optimal coil system with less stringent tolerances, such that small deviations cause reduced deterioration of the properties important for fusion performance.

Keywords: Stellarator, Nonlinear Coil Optimization, Reduced Engineering Tolerances, Robust Magnetic Field, Stochastic Optimization
Submitted to: *Nucl. Fusion*

1. Introduction

Stellarators confine plasma in a 'magnetic cage', which is produced by a set of external coils. These coil systems need to be optimized to fit both the physics requirements of the

corresponding vacuum field and the engineering requirements of the coil structure. The accuracy of the constructed coil system is crucial for the performance of the stellarator and is a cost and schedule driver as exemplified below.

The Wendelstein 7-X (W7-X) project required relative coil tolerances (defined as the allowed tolerance divided by the average coil radius) between 0.1% - 0.17% [1] . These strict tolerances were kept, and the device was put successfully in operation with a remarkable precision of its magnetic field topology [2]. However, precision requirements for the magnets were a major challenge and affected both cost and schedule negatively, as stated in [3]: "The assembly process which took about 1 000 000 man-hours up to March 2014, was essentially dominated by the high demands on tolerances for the position of the superconducting coils". The construction of the National Compact Stellarator Experiment (NCSX), which started in 2003, required relative coil deviation tolerances of about 0.08%. The associated difficulty and risk "was recognized but underappreciated at the project outset" [4]. Unfortunately, "the budget increases, schedule delays and continuing uncertainties of the NCSX construction project necessitate its closure" [5] in 2008.

For the design of the Columbia Nonneutral Torus (CNT), whose design started around 2001, well after the coil designs of W7-X and NCSX had started, the issue of tight coil tolerances was addressed from the beginning of the coil design. A Monte-Carlo-type perturbation analysis was performed as an integral part of determining the coil currents and coil locations [6]. A configuration was chosen that exhibited large flux surfaces even in the presence of minor coil displacements, over others that nominally had larger confinement volumes but were much more sensitive to coil placement inaccuracies. The chosen configuration had assembly tolerances of 0.5% - 1% for deviations in coil location and orientation [7], ie. about an order of magnitude looser than W7-X and NCSX. The resilience against error fields allowed the acceptance of several discrepancies. The robustness of the volume of the magnetic surfaces was directly verified experimentally [8], and significant coil displacements were later confirmed with state-of-the-art metrology [9], proving that the device was not accurately built but rather that its design point indeed was robust against coil manufacturing and installation errors.

Those results were promising, and very important for the successful construction of CNT, but the optimization goals for CNT were simple: to maximize the volume of good flux surfaces for a stellarator with only four circular coils. It was until now not clear if the approach taken would be effective for the much more complicated and multi-faceted optimization goals for stellarators designed for fusion energy research, and for the much more complicated coils that result from state-of-the-art coil design codes. The approach itself, described in more detail in [6], is one that optimizes a cloud of coil sets in the near vicinity of a particular configuration that is being optimized. This configuration will be referred to as the leading configuration and is in the center of the multidimensional cloud of other configurations, a Monte-Carlo sample of the coil configurations that could be realized within the uncertainties associated with construction and assembly of an actual device. The average performance of this cloud of coil sets is ascribed to the

leading configuration. By sampling neighboring states in this way, a penalty is given to configurations whose performance could be well below that of the leading configuration. In particular, this sample average avoids that a search ends up in a narrow "peak" (if maximizing) or trough (if minimizing) which is then, due to the narrowness, not robust against errors in the range that is to be expected. This technique also effectively smoothens the optimization landscape, as illustrated for a toy model in Figure 1. In this case, the non-robust optimum for $x \approx 0.42$ vanishes and the most stable optimum emerges as the global optimum at $x \approx 0.3$ in both locally averaged parameter spaces. For the one using a broad Gaussian distribution, the function to be optimized is now very smooth, and it even has just one maximum. For the one using a narrower Gaussian, the smoothening is also evident, and the global maximum is the robust one, but several other local maxima are still present.

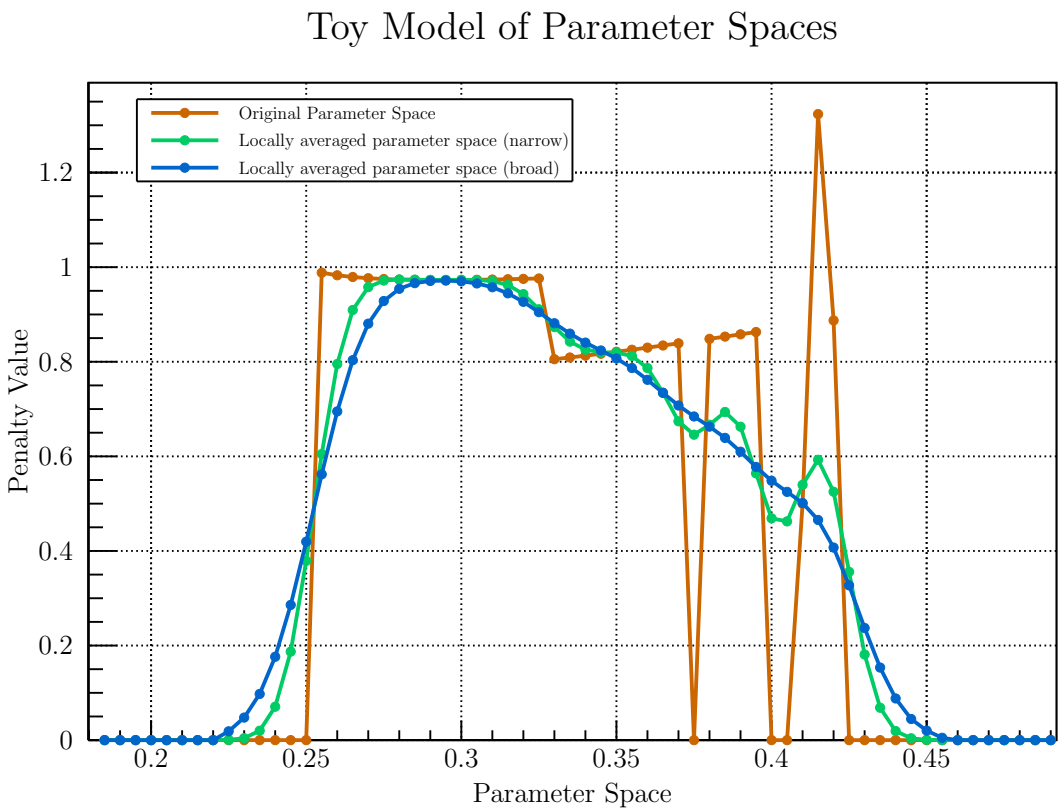

Figure 1: In stellarator optimization, the target function may be a non-smooth function of control parameters, and may exhibit a global maximum that is not robustly attainable since the control parameters (such as coil shapes) have finite tolerances. Taking the average of normal distributed samples in a small region of control parameters has a smoothening effect on the parameter space, and avoids getting stuck in non-robust optima. This is illustrated here for a toy model.

The approach presented in this paper follows the same philosophy but is applied for a much more fusion relevant problem addressed with state-of-the-art codes. Nonlinear coil optimization is expanded by an iterative perturbation analysis, which aims to

achieve higher engineering tolerances through the aforementioned technique. Instead of optimizing a single coil set, we optimize the sample average of $N+1$ coil sets, where N is the number of additional samples. Its disadvantage, that $N+1$ times more computational resources are needed, is partly offset by the highly efficient parallelization that is possible for this problem.

In Section 2, we transform the original optimization problem into a stochastic optimization problem. We then expand upon the sampling technique first performed for CNT, applying a more sophisticated version of it to a much more complicated optimization problem, namely that of W7-X. Since the original optimization sequence for the W7-X coil system is not available, a new coil optimization sequence was developed. Additionally, in order to achieve a high amount of experimental flexibility, the W7-X coil configuration was designed to produce multiple plasma configurations. The various features resulted in a compromise between different, at times conflicting, optimization targets. Consequently, the W7-X coil system does not correspond precisely to the original plasma boundary designed by J. Nührenberg [10] and hence, we do not use the present W7-X coil set as a reference case to validate how good our newly optimized coil sets are. Instead, we compare our results with a coil set that was found with our most recent standard optimization - one that only optimizes for a single coil set at a time without averaging over a cloud of perturbed coils. The test is described and analyzed in Section 3. We conclude in Section 4 and give an outlook on our future projects in Section 5. More details about the optimization sequence can be found in Appendix A.

2. Stochastic Optimization Problem

In nonlinear coil optimization one starts with an initial coil configuration $x \in X \subset \mathbb{R}^n$, where n is the number of parameters describing the coil set, and tries to solve the optimization problem

$$\min_{x \in X} f(x) \tag{1}$$

with a nonlinear optimizer. The objective function $f : X \to \mathbb{R}$ yields the fitness of the corresponding coil set x and is the measure of optimization. It summarizes the differences of the quality criteria $q_i(x)$ and their design value q_i^{design} in a sum of squares:

$$f(x) = \sum_{i=1}^{k} \omega_i \left(q_i(x) - q_i^{design} \right)^2 . \tag{2}$$

Here, the ω_i denote the weight constants. Since on the one hand the coil geometry and installation has a certain precision, and deviations within this precision in general are not predictable, and on the other hand our goal is to be able to relax precision requirements with minimal loss of plasma performance, the optimization problem is a stochastic one, with the following formulation:

$$\min_{x \in X} \{ F(x) := \mathbb{E}_{P_x}[f(\xi(x))] \} \qquad \xi(x) = x + \epsilon \quad \text{with} \quad |\epsilon| \ll |x|. \tag{3}$$

The expectation is taken with respect to the probability distribution P_x of the random vector ξ. This includes random deviations of a manufactured and installed coil set from its as-designed geometry and placement. We therefore optimize for uncertainty of the deformation of the coil set during construction. The probability distribution P_x carries the information of the likelihood of any possible geometric deviation of each particular coil set $x \in X$. P_x is assumed to follow a normal distribution.
We follow the approach taken in [11] and use a Monte Carlo sampling approach to approximate the expected value $\mathbb{E}_{P_x}[f(\xi(x))]$. Let $\xi^1(x), \ldots, \xi^N(x)$ be i.i.d. (independently and identically distributed) samples generated by the n-dimensional normal distribution $\mathcal{N}(x, \Sigma)$. Here, the mean of the normal distribution is our initial coil configuration $x \in X$ and the standard deviation is equal in every dimension and chosen appropriately to obtain the desired deviation in the euclidean space. Then by the Law of Large Numbers, for a given $x \in X$, we have

$$F_N(x) := (N+1)^{-1} \sum_{i=0}^{N} f(\xi^i(x)) \longrightarrow \mathbb{E}_{P_x}[f(\xi(x))], \text{ for } \; N \to \infty \tag{4}$$

almost surely, where $\xi^0(x) := x$ is the unperturbed coil set (leading configuration). The sample average F_N is an unbiased and consistent estimator of $F(x) = \mathbb{E}_{P_x}[f(\xi(x))]$. The convergence in probability is of order $O_p(N^{-\frac{1}{2}})$, implying that a ten-fold increase in accuracy of the estimate of the expectation requires a 100-fold increase in sample size.

3. Results

We test our stochastic optimization technique using a newly developed coil optimization sequence for the W7-X plasma boundary [12]. The sequence consists of six optimization runs (cf. Table A1) such that the previous optimization run is used as a new starting point. The initial coil set is computed with NESCOIL [13] on a winding surface (WS) located 45 cm beyond the outer plasma surface. The objective function f, which assigns a penalty value to every coil set, is calculated by ONSET [14]. The evaluation criteria used in ONSET are shown in Table A1. A coil is defined by 12 periodic spline points and raised to 13 spline points in optimization run 4. Each spline point is defined by a poloidal angle u and a toroidal angle v. The 5-fold symmetry of W7-X combined with the stellarator symmetry reduces the primary coil set to 5 different coils. The WS is interpolated with 13 parameters between two limiting surfaces and since the auxiliary coils are not being considered we arrive at $n = 133$ or rather $n = 143$ parameters in total. We have chosen Brent's PRincipal AXIS algorithm [15] for the nonlinear optimization without using derivatives.

We modified the objective function from calculating the penalty value $f(x)$ of a single coil set $x \in X$ to generating a cloud of coil sets $\xi^1(x), \ldots, \xi^N(x)$, which are slight variations of the parameters $x \in X$, computing the corresponding penalty values $f(\xi^1(x)), \ldots, f(\xi^N(x))$ and assigning the average value $F_N(x)$ to the unperturbed leading

configuration. A deformation of the parameters used for the optimization results in coil deformations that preserve the 5-fold and stellarator symmetry of W7-X and target systematic errors that are equal for each coil of the same type during coil fabrication. The way this stochastic optimization problem is approached, one must provide as input a relevant scale length over which the coil deviations are distributed. We determined this as follows: During fabrication of the W7-X winding pack, the average deviation of the non-planar coils from their Computer-Aided-Design (CAD) shape were approximately 2 mm [16]. In the subsequent assembly, reference marks were used to guide and validate the assembly. "In total no reference mark co-ordinate deviated more than 5.7 mm from its manufacture value" [17]. We take these values to define the rough length scale over which we numerically displace coils in our stochastic optimization, since this was actually achieved (and therefore achievable), but only achieved with a significant engineering and metrology effort [3]. Thus, we chose the n-dimensional normal distribution $\mathcal{N}(x, \Sigma)$ appropriately to optimize numerically for robustness against deviations that are 2 mm on average, with greater than 6 mm deviations being present in roughly 10% of the numerically tracked geometric points on the coils in the Euclidean space.

Each time the algorithm computes the modified objective function $F_N(x)$ and evaluates the fitness of the coil set x, the perturbed coil sets $\xi^1(x), \ldots, \xi^N(x)$ are defined anew. Thus, the deviations are not only randomly assigned once at the beginning of every optimization sequence but continue to be re-randomized for each evaluation of $F_N(x)$. To provide a comparison basis, the optimization sequence was designed in the classical way $F_1 = f$ with $N = 1$ samples, where only one coil set is evaluated in each optimization step. We call this the reference coil set with cloud size of 1. We then increased the sample size N and optimized with a cloud size of 100, 1000, 2000, 4000, and 8000 perturbed coil sets, so as to determine if there is an optimum cloud size, and some convergence or saturation of the results at very large cloud size. We restricted the coil displacements to variations of the parameters that move the circular spline (coil) along the two-dimensional WS, instead of allowing coil deviations also perpendicular to this surface. This restriction was done in order to keep the computational effort moderate for this first study and we expect to be able to lift this restriction in future studies.

We compare the results of the first optimization run in Section 3.1 by comparing the penalty values. In Section 3.2 we compare the results after the optimization sequence has completed, ie. it has come to a stationary value of F_N. The robustness of the newly optimized coil sets is assessed by plotting the penalty value distributions. They are obtained by deviating the coil sets multiple times at the aforementioned level (average deviation 2 mm). A narrower histogram gives more confidence in the performance, ie. it is more robust against coil deviations of the size assumed here. Another important question arises if the histogram not only narrows (more robust), but also shifts to the left (lower penalty values) indicating that a better optimum has been found, and that the device would accordingly perform better than the device optimized with the classical algorithm. Furthermore, the robustness can also be quantified in terms of relaxed

tolerances, for a given performance. For that, we compare the classically optimized coil set to the best coil set obtained with the new optimization technique.

3.1. Comparison after the 1st optimization run

In Table 1 we compare the results of the 1st optimization run quantified by the penalty value. Since the algorithm is non-deterministic, we assess the results by running the 1st optimization multiple times, as suggested in [11]. Each dot in Figure 2 represents

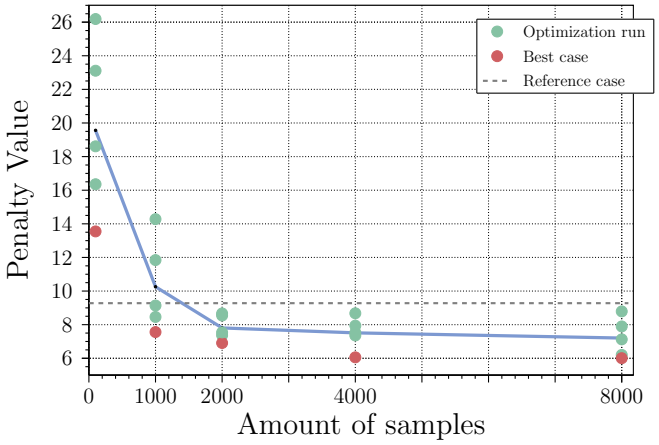

Samples	#	Best	∅	σ
1	-	9.28	-	-
100	5	13.55	19.56	4.56
1000	5	7.56	10.26	2.47
2000	5	6.91	7.81	0.69
4000	5	6.06	7.51	0.86
8000	5	6.00	7.21	0.95

Figure 2 & Table 1: Results of the 1st optimization run are listed in form of the best coil set (Best), the average penalty value (∅) and the standard deviation (σ)

an independent 1st optimization run. The average penalty value is shown in the blue line. It declines the more samples are used, as does the best result (red dots & Best column). Using the stochastic optimization with a sample size less than or equal to 1000, we observe that the average penalty value is worse than the reference optimization. It indicates that in the transformed parameter space more local optima are present when the sample size is low than in the parameter space of the reference optimization, which causes the optimizer to halt even earlier. Increasing the sample size leads to the intersection between the blue and dashed line, where both parameter spaces have a similar landscape such that both optimization techniques show the same performance. Increasing the sample size further intensifies the smoothing effect and leads to a reduction of local optima which causes the optimizer to find even better results than the reference optimization.
The standard deviation shows that the statistical noise of the penalty values after the 1st optimization run is very high for low sample rates, in addition to the bad performance of the coil sets. A plausible explanation of both negative behaviours is the combination of random samples and the discrepancy between the dimension of the parameter space ($n = 133$) and a sample size of the same magnitude. In each evaluation of the penalty function, the chance is high that not every dimension is covered and the optimal path in

the vicinity of the leading configuration is not visible for the optimizer. The stochastic behaviour of our sampling technique guarantees that these situation happen, sometimes more, sometimes less. In addition to the bad convergence rate of the expected value mentioned in Section 2, this effect leads to an overall poor approximation of the expected value when the sample size is low.
Unfortunately, we do not fully understand all of the effects and their interplay which cause the parameter space to contain more local optima than the reference case for low sample rates but we we intend to include a more detailed investigation in our future studies. In general, all statistical measures converge as the sample size increases.

3.2. Comparison after completion of the optimization sequence

We compare the robustness of the design of the coil sets after their last optimization run (cf. Table A1) by analyzing the penalty value distribution for clouds of 100 000 perturbed coil sets around the newly optimized leading coil sets. The entries of 100 000 provide enough statistics to compute the relevant variables of the histogram. The results are shown in Figure 3. The coil shapes are parametrized such that the 2 mm average deviation in the Euclidean is not automatically kept, even if the coil shape parameter deviation is kept fixed. Since the WS is part of the optimization and changes its form, the perturbation in the Euclidean space slightly increased over the course of the optimization sequence in the cases with 100, 1000 and 2000 samples. The perturbations in Figure 3 are done with the same technique as described in Section 2 but normalized to reach again an average deviation of 2 mm in the Euclidean space. The coil sets are compared by their penalty value $f(x_0)$ and their robustness, which is seen by an increased height and decreased width of the distribution function. Additionally, the high-end tails of the penalty function distributions are analyzed.

In Figure 3 the leading configuration of the optimization with 8000 samples reaches the best result of $f(x_0) \approx 5.39$, and it is nearly identical to the result for a cloud size of 4000, indicating a possible convergence of the penalty values. The reference coil set (using the classical, non-stochastic, optimization) has $f(x_0) \approx 6.65$. Thus, we observe a systematic improvement of this nonlinear coil optimization - irrespective of the issue of robustness, and we see that we need a cloud size above 2000 before the new algorithm outperforms the old one. The width of the distribution is improved already for the smallest cloud size (100), but then actually begins to get wider again for the largest cloud sizes - the coil set optimized with a 2000 sample cloud reaches the highest peak and together with the case of 1000 samples they reach the smallest widths (RMS) of the distribution function. Therefore, the case with 2000 samples is more "robust" in the sense of stable quality criteria, but will have poorer performance than the 8000 sample optimization. Overall, both robustness and average performance are significantly improved over the reference coil set. The double peak of the histogram of the reference coil set is not understood in any detail at this point, but will be further investigated in the future.

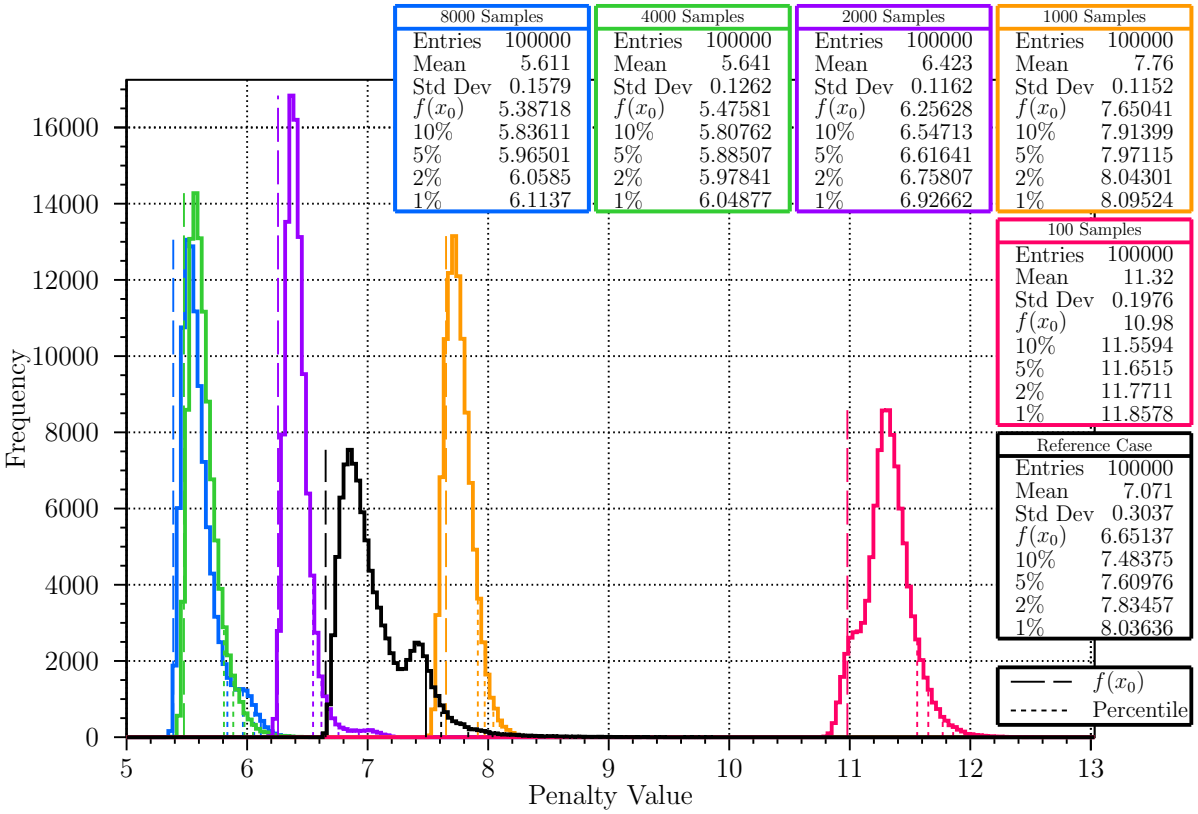

Figure 3: The result of all optimizations including the reference coil set (black) are randomly perturbed 100 000 times and the penalty value distribution is displayed.

The results in Figure 3 are positive, despite this slight broadening of the distribution function for the 4000 and 8000 sample sets relative to the results obtained with the sample size of 2000. From an experimental/user perspective, what is desired is confidence in good performance at tolerances that are relaxed enough that they do not affect the cost or schedule of construction significantly. It is clear from Figure 3 that this optimization is a significant win at the 2 mm accuracy level. This can be quantified further, and put in relation to relaxed engineering tolerances by focusing on the high-end tail of the histogram - the poorly performing configurations.

We plot in Figure 4 the development of the percentiles 10%, 5%, 2% and 1% of the penalty function histogram, thereby quantifying what performance is guaranteed at the 90%, 95%, 98% and 99% confidence level. We compare these percentiles of the reference coil set (classical optimization) and the optimization with 8000 samples (our most comprehensive stochastic optimization). For the 2 mm accuracy case which was used for the optimization and achieved in W7-X (left side of graph) it is again clear, as evident in Figure 3, that the new configuration will outperform the old with better than 99% likelihood. In Figure 4 we now increase the coil tolerances progressively (x-axis) and monitor how the penalty value for these percentiles increases (y-axis). At the 99% confidence level, the classical optimization would have led to a penalty value up to about 8, given the 2 mm tolerances achieved in W7-X. For the new optimization,

the same would be achieved with 99% confidence with coil tolerances of approximately 4.5 mm, more than a factor of two relaxation for the same risk taken. If we lower the degree of confidence (90%, 95%, 98% confidence) the relaxation of tolerances is a bit larger, approaching a factor of three, eg. 5.8 mm tolerances at the 90% confidence level. These comparison points are given as stars. The very steep ascents beyond these values are possibly due to the combination of low statistics and discontinuities in the penalty function, but should be taken with some caution at this point. Regardless of this issue, the data fully supports the statement that the stochastic optimization would give significantly better performance at the same level of engineering tolerances (Fig 3), and comparable if not better performance with a factor of two relaxed tolerances, for a variety of confidence standards.

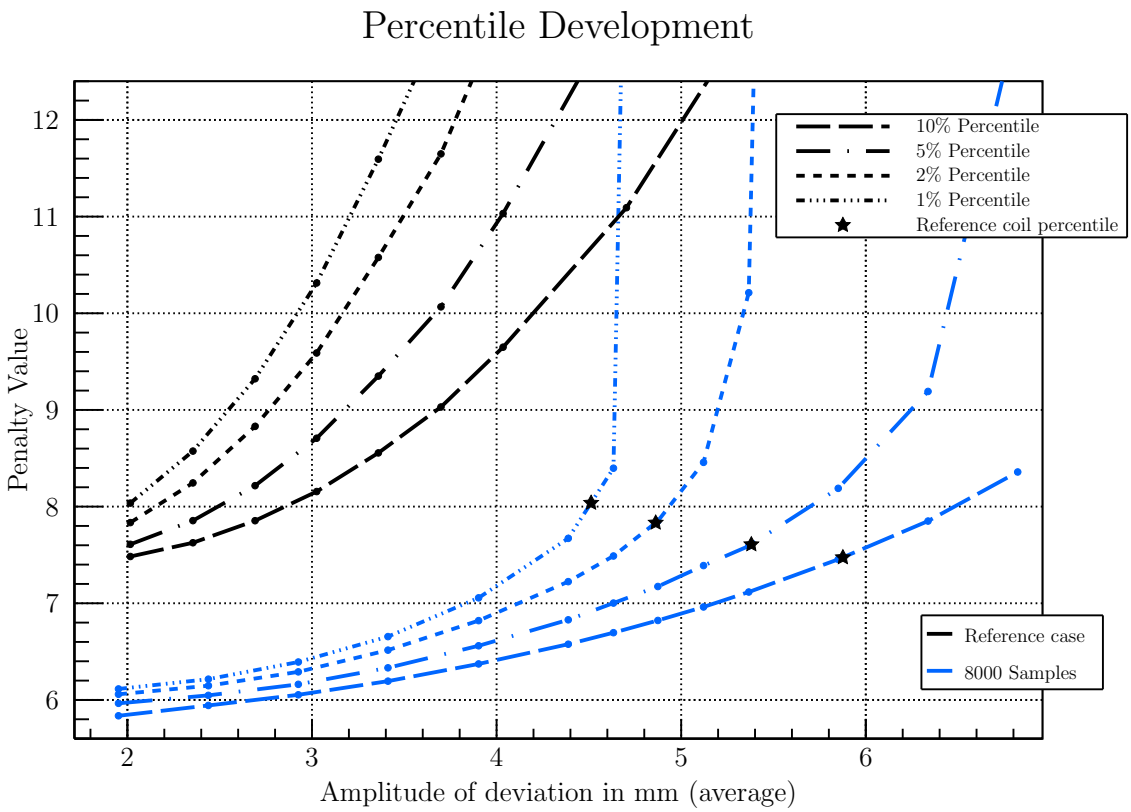

Figure 4: Percentile development of the reference coil case and the optimization with 8000 samples.

3.3. Remark

The object of optimization in this study is the approximation of the expected value which aims to represent the most likely realization of a stellarator. The relaxation of precision requirements is no direct target of the optimization since it is no criterion in the objective function. It is only indirectly represented in the sample average F_N but nevertheless leads to a coil configuration with significantly relaxed tolerances.

The optimization suite ONSET is capable of optimizing the three-dimensional shape of the WS within two constraining toroidal surfaces. This offers the possibility to extend the coil displacements from a non-planar two-dimensional in-surface deviation to also capture three-dimensional deviations. A first attempt at using this feature, which demands significantly more computer resources, has been made and will be reported on in a future publication. Initial results appear consistent with what has been reported here, but convergence studies have not yet been completed.

4. Discussion and conclusion

A stochastic technique of nonlinear coil optimization for stellarators that includes engineering tolerances has been developed and tested on a state-of-the-art optimization problem. The technique shows significant promise. Improved performance is seen - a 20% decrease in the targeted penalty function. The new optimization shows more engineering robustness - the results are less sensitive to coil displacements, and one can relax coil tolerances by at least a factor of two relative to a coil design created using earlier techniques. The relaxation of engineering tolerances was expected. But it was not necessarily expected that the algorithm found a better optimum - irrespective of engineering tolerances. As illustrated in Figure 1, this technique can have a smoothing effect in the otherwise rather spiky optimization space of stellarators, and this may have allowed the algorithm to find a better global optimum than that found by the earlier algorithm. Improved performance and then convergence with sample (cloud) size was seen, with the results from cloud sizes of 4000 and 8000 being nearly identical.

5. Outlook

This work brings up several questions, some of which can be addressed soon. Coil deviations in all three dimensions (not just within a prescribed toroidal surface, and not just stellarator-symmetric) will be addressed next. Additionally, a penalty on the width of the distribution will be included in the objective function in order to directly target reduced precision requirements. The reason why the original algorithm got stuck in a non-global minimum together with the origin of the double peak in the histogram of the reference coil set will be investigated. The new, more optimal, configuration will be compared to the configuration found with the earlier algorithm, in particular with respect to each individual physics target lumped into the penalty function, and the coil geometry will be investigated. At first glance, the geometry of the coils is somewhat different for the newly found coil set but does not appear to be more complicated than what was found with the standard algorithm.

6. Pictures

We compare the shape of the coil set of the optimization with 8000 samples in blue with the reference coil optimization in black in Figure 5 & 6. We show 2 perspectives of one stellarator segment where two half modules are connected at the triangular plane. Figure 5 visualizes the differences from the inside of the torus towards the outside and Figure 6 shows the opposite direction from the outside towards the inside. In general, the two coil sets are quite similar with a certain tendency that the blue coil set is slightly less windy than the reference coil case in black.

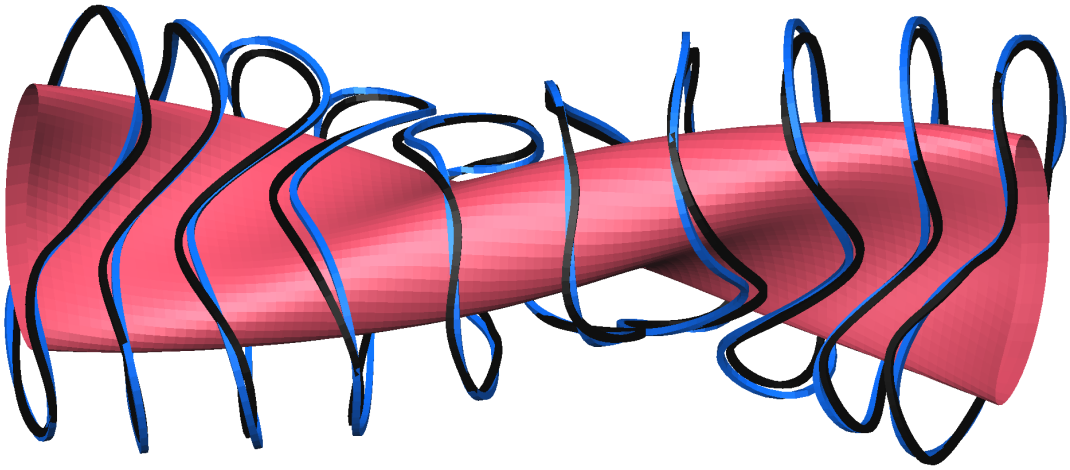

Figure 5: View of two half modules connected at the triangular plane from the inside towards the outside. The coil set in blue is the optimization with 8000 samples and the coil set in black is the reference coil optimization.

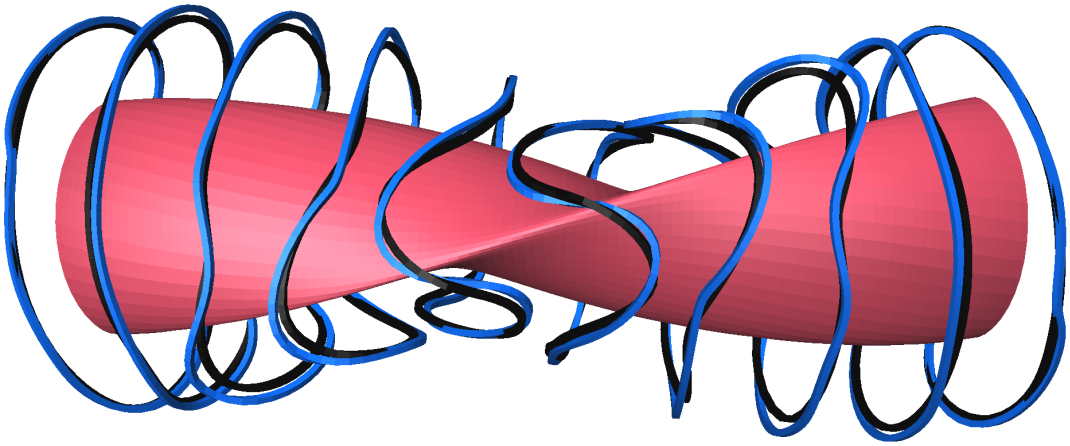

Figure 6: View of two half modules connected at the triangular plane from the outside towards the inside. The coil set in blue is the optimization with 8000 samples and the coil set in black is the reference coil optimization.

7. Acknowledgement

This work has been carried out within the framework of the EUROfusion Consortium and has received funding from the Euratom research and training programme 2014-2018 under grant agreement No 633053. The views and opinions expressed herein do not necessarily reflect those of the European Commission.

References

[1] Th. Rummel et al., Accuracy of the construction of the superconducting coils for Wendelstein 7-X, *IEEE Trans. Appl. Supercond.* 14(2004)(2(June)), pp. 1394-1398

[2] T. Sunn Pedersen et al., Confirmation of the topology of the Wendelstein 7-X magnetic field to better than 1:100,000, *Nature Communications* (2016) Vol. 7, 13493

[3] H.-S. Bosch et al., Final integration, commissioning and start of the Wendelstein 7-X stellarator operation, *Nuclear Fusion*, Vol. 57, 116015 (2017)

[4] G.H. Neilson, Lessons Learned in Risk Management on NCSX, *IEEE Transactions on Plasma Science*, (2010) Vol. 38, No. 3 pp. 320-327

[5] R.L. Orbach, Statement about the Future of the Princeton Plasma Physics Laboratory, *Under Secretary for Science and Director, Office of Science*, U.S. Department of Energy, May 22, 2008

[6] Jason Paul Kremer (2006) *The Creation and First Studies of Electron Plasmas in the Columbia Non-neutral Torus* (Doctoral dissertation) Retrieved from ProQuest Dissertations and Theses. (Accession Order No. AAI3249101)

[7] T. Sunn Pedersen et al., Construction and initial operation of the Columbia Non-neutral Torus, *Fusion Science and Technology*, Vol. 50, (2006) pp. 372

[8] T. Sunn Pedersen et al., Experimental demonstration of a compact stellarator magnetic trap using four circular coils, *Physics of Plasmas* **13**, 012502 (2006)

[9] K.C. Hammond et al., Experimental and numerical study of error fields in the CNT stellarator, *Plasma Physics and Controlled Fusion*, Vol. 58, No. 7 (2016) pp. 074002

[10] Nührenberg J. and Zille R. 1988 *Physics Letters* A **129** 113-117

[11] A. Shapiro, Stochastic programming by monte carlo simulation methods, *ESAIM: PROCEEDINGS*, December (2003), Vol. 13 pp. 65-73

[12] J. Lobsien, M. Drevlak, Status of Coil Optimisation with ONSET, *Manuscript in preparation.*

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

[14] M. Drevlak, 20th Symposium on Fusion Technology, Marseille, France (1998) pp. 883.

[15] K. Gegenfurtner, PRAXIS: Brent's algorithm for function minimization, *Behavior Research Methods, Instruments, & Computers*, (1992) Vol. 24, No. 4 pp. 560-564

[16] T. Andreeva et al., Influence of construction errors on Wendelstein 7-X magnetic configurations, *Fusion Engineering and Design*, 84(2)(2009), p.408 - 412

[17] T. Andreeva et al., Tracking of the magnet system geometry during Wendelstein 7-X construction to achieve the designed magnetic field, *Nucl. Fusion*, 55(6)(2015), p.063025.

[18] Michael Drevlak, Automated Optimization of Stellarator Coils, *Fusion Technology*, Vol. 33, No. 2, Page 106-117, Year 1998

Appendices

Appendix A. Details about the optimization sequence

The optimization sequence is listed in Table A1, which illustrates the change of the weight constants throughout the six optimization runs. Every entry consists of two values (design value/ weight constant) which define the objective function f. Exceptions are the limiting surfaces (inner surface / outer surface), and the number of spline points which are defined per coil. The optimization sequence is divided into the optimization of the field error under observance of the geometric restrictions (1-4) followed by the optimization of the properties of the magnetic field. In the final step we compute additionally the Fourier coefficients of a magnetic surface inside the plasma boundary. The difference to the Fourier coefficients obtained from the target magnetic surface inside the original plasma boundary designed by Nührenberg [10] is then minimized in the last optimization run. The optimization of the field error took 4 runs. We refer to [12] for a detailed description of the optimization sequence and the design criteria.

Table A1: Optimization Sequence

Optimization run:	**1**	**2**	**3**	**4**	**5**	**6**
Maximum Field Error	$0/\ 1.0 \cdot 10^3$	·	·	·	·	·
Mean Field Error	$0/\ 1.0 \cdot 10^5$	·	·	·	·	·
Clearance	0.35/1	0.3/1	0.27/1	·	·	·
Curvature	3/0.7	·	·	·	·	·
Distortion	0.3/0.4	·	·	·	·	·
Magnetic Axis (bean)	-	-	-	-	5.93/4	·
Magnetic Axis (triangle)	-	-	-	-	5.17/15	5.17/450
Magnetic Ripple on Axis	-	-	-	-	0.11/3	·
Iota on Axis	-	-	-	-	$0.88/1.0 \cdot 10^3$	·
Magnetic Shear	-	-	-	-	1.56/0.4	1.56/0.8
Magnetic Well	-	-	-	-	$0.007/1.0 \cdot 10^3$	·
Fourier Coeff. of inner Surface	-	-	-	-	-	varies/$1.0 \cdot 10^2$
Limiting Surfaces	+30/+60	·	+30/+65	+30/+70	·	·
Points per Coil	12	·	·	13	·	·
Points of WS	13	·	·	·	·	·

Description:

- Target not included in objective function.

· Value from the previous run was used.

Field Error Normal magnetic field on the plasma boundary.

Clearance Minimum distance between adjacent coils.

Curvature Maximum coil curvature $\frac{d^2x}{ds^2}$.

Distortion Weighted curvature defined in [18].

Magnetic Axis Major radius.