

LETTER

Constrained stellarator coil curvature optimization with FOCUS

Thomas G. Kruger^{1,†}, C. Zhu², A. Bader¹, D. T. Anderson¹
and L. Singh¹

¹HSX Laboratory, University of Wisconsin, Madison, WI 53706, USA

²Princeton Plasma Physics Laboratory, Princeton University, Princeton, NJ 08540, USA

(Received 30 September 2020; revised 21 January 2021; accepted 25 January 2021)

Finding less complicated coils that have adequately low field errors is a crucial step in stellarator development. One coil metric that is of high importance is the maximum curvature of the coil centreline, or coil single filament. Conductors cannot be bent below some threshold minimum radius of curvature. High coil curvatures can cause strains to exceed acceptable levels, especially in superconducting coils. We investigate three ways to optimize coil curvature and find that applying penalty functions to the coil curvature solves for coils that have a constrained maximum curvature and low field error. Penalty functions are implemented in FOCUS and coil solutions optimized for an HSX-like ‘plasma boundary’ are presented.

Key words: plasma devices

1. Introduction

Stellarators possess inherent advantages that are important to consider for future fusion power plant designs, including steady-state operation, resiliency to disruptions and theoretically high beta stability limits (Gates *et al.* 2018). Although stellarators have desirable physics properties, they also require complicated engineering. Generally, stellarators do not require current drive to achieve reasonable rotational transform profiles. To produce non-axisymmetric magnetic fields, non-planar coils are typically used (Merkel 1987; Drevlak 1999; Landreman 2017). Non-planar coils complicate coil design, engineering, assembly and supports. Difficulties arising from non-planar coils led to significant delays during the construction of the W7-X experiment (Riße 2009) and to the cancellation of the NCSX experiment at the Princeton Plasma Physics Laboratory (Neilson *et al.* 2010).

There is a trade-off between coils that accurately reproduce a desired magnetic field and coils that are simple to manufacture. Stellarator coils need to satisfy engineering constraints such as maximum coil curvature, minimum coil-to-coil separation and minimum coil-to-plasma

† Email address for correspondence: tkruger@wisc.edu

separation. After such engineering constraints are satisfied one can also optimize for additional considerations such as coil resiliency to displacements, coil spacing for auxiliary systems, coil cost and coil complexity. Coil complexity is left ambiguous here and we do not intend to define a set of quantitative metrics for it. Coil complexity is necessarily dependent on the manufacturing process, whereas the coil maximum curvature is important independent of the manufacturing process. This letter investigates the maximum curvature engineering constraint, just one of many coil considerations.

Previously, both the ONSET and COILOPT codes used curvature objective functions to target coil curvature. The ONSET code uses a penalty function that quadratically penalizes curvatures that exceed some input maximum curvature (see (5) and (6) of Drevlak (1998)). This curvature objective function is not differentiable, which limits the applicability of the objective function to different optimization schemes. Additionally the curvature objective function is dependent on coil lengths. The COILOPT code used an exponentially weighted function to heavily penalize high coil curvatures (see (4) of Strickler, Berry & Hirshman (2002)). One difference between the COILOPT exponential function and our penalty function is that the exponential function penalizes all values of curvature, not just values that are above the maximum. Furthermore, in COILOPT, coils are constrained to lie on a predefined winding surface which greatly limits flexibility.

In this letter we compare three curvature objective functions. In doing so we attempt to elucidate the salient properties of penalty functions that constrain curvatures below some maximum value and objective functions that try to reduce the integrated curvature. While both the ONSET and COILOPT formulations target curvature and are able to reduce coil curvatures, there has not yet been a systematic effort to understand what specific choices of the curvature function yield. This letter provides a first attempt at resolving these questions.

The first two curvature objective functions minimize a coil's average curvature and average curvature squared. The third curvature objective function uses penalty functions to explicitly constrain a coil's maximum curvature. The maximum curvature constraint is of crucial importance since it directly determines whether a coil set can be constructed. For instance, if the minimum radius of curvature is greater than half the small dimension of a coil width, then the coil will self-intersect. Also, smaller radii of curvature directly increase coil strain and can increase the magnetic field strength in the coil, both of which can violate engineering constraints for superconducting coils.

To optimize for a coil's maximum curvature, we will use the single-filament coil optimization code FOCUS (Zhu *et al.* 2017). Single-filament coils are initial approximations where the coil cross-section is taken to be a point. This coil approximation is appropriate for any type of coil such as low-temperature superconducting, high-temperature superconducting and copper coils. Depending on the size of the coil cross-sections relative to the coil-plasma distances, a multi-filament coil optimization can be done as a subsequent step (Singh *et al.* 2020). For this letter it is assumed that coil means single-filament coil.

Our motivation and method for constraining maximum coil curvature are described in the rest of this letter. Newly implemented FOCUS objective functions are described in § 2. Coils optimized with these objective functions are shown in § 3. The objective functions and resulting coils are discussed in § 4. Conclusions and future work are discussed in § 5.

2. FOCUS objective functions

FOCUS is a stellarator coil optimization code that takes advantage of parallel computing. FOCUS solves nonlinear coil optimization problems of the following form:

$$\min \sum_i w_i f_i, \quad (2.1)$$

where w_i is a user-specified weighting associated with the objective function f_i . These objective functions are dependent on and minimized by changing coil geometries and currents. Objective functions in FOCUS are twice differentiable. These objective functions are minimized with gradient-based optimization algorithms, like the conjugate gradient method, or with Hessian-based algorithms, like the Newton method (Zhu *et al.* 2018). Coils in FOCUS are approximated as space curves and defined using truncated Fourier series as

$$\mathbf{r}(t) = x(t)\hat{\mathbf{x}} + y(t)\hat{\mathbf{y}} + z(t)\hat{\mathbf{z}} = \sum_{n=0}^{N_F} [(X_{c,n} \cos(nt) + X_{s,n} \sin(nt))\hat{\mathbf{x}} + (Y_{c,n} \cos(nt) + Y_{s,n} \sin(nt))\hat{\mathbf{y}} + (Z_{c,n} \cos(nt) + Z_{s,n} \sin(nt))\hat{\mathbf{z}}], \quad (2.2)$$

where $t \in [0, 2\pi]$ is our parameterizing variable and N_F defines the highest Fourier mode frequency to be used in the optimization. The optimization variables are the Fourier mode amplitudes $X_{c,n}$, $X_{s,n}$, $Y_{c,n}$, $Y_{s,n}$, $Z_{c,n}$ and $Z_{s,n}$ and coil current for each coil.

The primary objective of coil optimization is to satisfy a normal field boundary condition on an input surface. This surface can be arbitrary, but typically is a ‘plasma boundary’ from a fixed boundary equilibrium. The normal field boundary condition can also be arbitrary, but is typically set by plasma currents and any fixed external current sources. For this letter we use a quadratic, normalized flux objective function:

$$f_{Bn} = \int_S \frac{1}{2} \left(\frac{\mathbf{B}_v \cdot \mathbf{n} - T_{Bn}}{|\mathbf{B}_v|} \right)^2 dS, \quad (2.3)$$

where \mathbf{B}_v is the magnetic field from coils, \mathbf{n} is the input surface unit normal and T_{Bn} is the normal field boundary condition. This objective function will be referred to as a field error objective function. We optimize coils to reconstruct magnetic fields from a zero-beta, fixed boundary equilibrium. This allows us to use the plasma boundary from the equilibrium as our input surface and allows us to set the normal field boundary condition to zero.

If the Neumann boundary condition is set to zero, coil currents are not held fixed and the quadratic flux objective function is not normalized by field strength, an additional objective function that optimizes for toroidal flux needs to be included. If it is not included in the optimization, a trivial solution will be solved for where coil currents go to zero, causing the normal field on the boundary to consequently also be zero. For this letter we keep coil currents held fixed, and use the normalized flux objective function (2.3); therefore a toroidal flux objective function is not needed.

‘Coil ripple’ errors are caused because field strength is large at surface positions close to coils and smaller in between coils. These errors are associated with high toroidal mode numbers corresponding to the number of coils. Coil ripple can be reduced by moving coils farther away from the boundary. Coil ripple errors are visible as varying stripes of minimum and maximum error on the boundary, as in figure 1(a). At larger coil distances it is more difficult to match the general plasma boundary shape even if the field is generated by an external current-carrying surface (Landreman & Boozer 2016). We refer to errors of this type as ‘non-ripple errors’.

The inverse of the Biot–Savart law is ill-posed and additional regularization is required in coil optimizations. FOCUS uses coil length as one of the primary regularizations. If there is no penalty on coil length, coils can become arbitrarily long to reduce coil ripple.

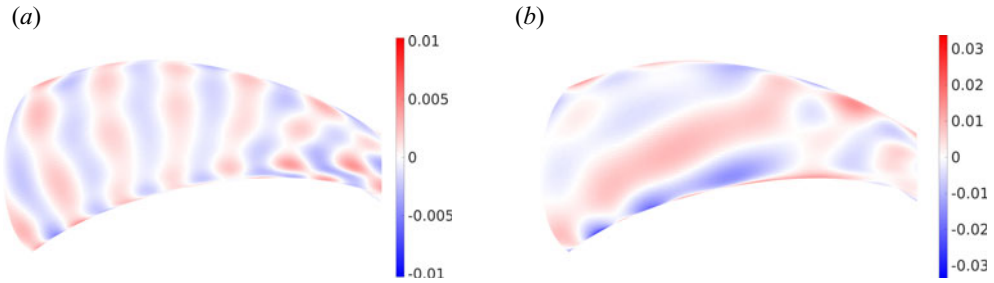


FIGURE 1. The $(\mathbf{B}_v \cdot \mathbf{n})/|\mathbf{B}_v|$ distributions plotted on a half-period of the HSX boundary from coils optimized with our penalty curvature objective function (a) and from coils optimized with a quadratic curvature objective function (b).

To obtain the results in this letter we use a quadratic penalty on coil length given as

$$f_L = \frac{1}{N_c} \sum_{i=1}^{N_c} \frac{1}{2} \frac{(L_i - L_{i,0})^2}{L_{i,0}^2}, \quad (2.4)$$

where N_c is the total number of coils, L_i is the length of the i th coil and $L_{i,0}$ is the user-specified target length of the i th coil. Including a curvature objective function further regularizes the optimization, but it is still important to include a coil length objective function to ensure that coil solutions are reasonable.

Coils optimized with only field error and length objective functions are discussed in § 3. These coils have a maximum curvature that is roughly an order of magnitude larger than that of the Helically Symmetric Experiment (HSX) coils, and therefore motivate the need for a curvature objective function. We now present three candidate curvature objective functions. Including a curvature objective function gives an optimization problem of the following form:

$$\min w_{Bn} f_{Bn} + w_L f_L + w_\kappa f_{\kappa,j}, \quad (2.5)$$

where $j = 1, 2, 3$ correspond to three possible forms of the curvature objective function.

The first curvature objective function ($j = 1$) is given as

$$f_{\kappa,1} = \frac{1}{N_c} \sum_{i=1}^{N_c} \frac{1}{L_i} \int_0^{2\pi} \kappa_i |\mathbf{r}'_i| \, dt, \quad (2.6)$$

$$\kappa = \frac{|\mathbf{r}' \times \mathbf{r}''|}{|\mathbf{r}'|^3}, \quad (2.7)$$

with each coil's curvature, κ_i , appearing linearly. Alternatively, it may be desirable to penalize areas of high curvature more than areas of low curvature. A second curvature objective function that optimizes for the square of the curvature should accomplish this. This function ($j = 2$) is given as

$$f_{\kappa,2} = \frac{1}{N_c} \sum_{i=1}^{N_c} \frac{1}{L_i} \int_0^{2\pi} \kappa_i^2 |\mathbf{r}'_i| \, dt. \quad (2.8)$$

It is important that we integrate over the coil's arc length $|\mathbf{r}'_i| \, dt$ and not the parameterizing variable dt . If we were to integrate over the parameterizing variable, the objective function




FIGURE 2. (a) The penalty function used in the $f_{\kappa,3}$ objective function, where the variable κ_i is changed to x and the value of κ_0 is set to 5. Multiple values of the penalty variable α are plotted. (b) The derivative of this penalty function.

would be parameterization-dependent. This means coils with the same shape could have different parameterizations and consequently different values for the objective functions. It is also important that we normalize each integral by the corresponding coil length. If we do not do this, coil lengths can decrease to minimize the objective functions despite having regions of higher curvature. The curvature objective functions (2.6) and (2.8) attempt to minimize an integrated function of coil curvature, but they do not constrain it. Therefore we are solving an unconstrained optimization problem. In order to constrain the coil curvature, the third curvature objective function ($j = 3$) uses a penalty function, making this a constrained optimization problem.

We apply a constraint on coil curvature so that it is below some maximum value κ_0 at all points in the set of coils. That is, $\kappa_i(t) \leq \kappa_0 \forall t$ and $\forall i$, where R_0 corresponds to the minimum radius of curvature. To apply this constraint we can use the following third curvature objective function which uses a penalty function:

$$f_{\kappa,3} = \left. \begin{aligned} & \frac{1}{N_c} \sum_{i=1}^{N_c} \frac{1}{L_i} \int_0^{2\pi} H_{\kappa_0}(\kappa_i) (\cosh(\alpha(\kappa_i - \kappa_0)) - 1)^2 |r'_i| dt, \\ & \alpha > 0 \quad \kappa_0 > 0, \end{aligned} \right\} \quad (2.9)$$

$$H_{\kappa_0}(\kappa_i) \equiv H(\kappa_i - \kappa_0) = \begin{cases} 0, & \kappa_i < \kappa_0, \\ \frac{1}{2}, & \kappa_i = \kappa_0, \\ 1, & \kappa_i > \kappa_0. \end{cases} \quad (2.10)$$

The third curvature objective function includes a Heaviside function defined in (2.10) which is discontinuous. It is important to notice that at the Heaviside function discontinuity, the term $\cosh(\alpha(\kappa_i - \kappa_0)) - 1$ goes to zero which can be seen in figure 2. This preserves the objective function’s continuity and moreover its differentiability. This objective function is twice differentiable, $f_{\kappa,3} \in C^2$. FOCUS has multiple optimization algorithms implemented. If the user were to choose one of the quasi-Newton methods, superlinear convergence can be observed since the FOCUS objective functions are twice differentiable (Nocedal & Wright 2006).

The penalty function only penalizes coil locations where coil curvature exceeds the specified maximum curvature, κ_0 . The value of $H_{\kappa_0}(\kappa_i)(\cosh(\alpha(\kappa_i - \kappa_0)) - 1)^2$ is zero if $\kappa_i \leq \kappa_0$, positive if $\kappa_i > \kappa_0$ and increases exponentially with the constraint violation $\kappa_i - \kappa_0$. To guarantee a constraint on coil curvature, the value of α needs to be iteratively increased to ∞ . Here α is called our penalty variable. As $\alpha \rightarrow \infty$, the value of $f_{\kappa,3} \rightarrow \infty$, if any coil has a location where the coil’s curvature exceeds the maximum curvature.

Optimization run	w_{Bn}	w_L	w_κ
No curvature objective function	1×10^4	1×10^3	0
Linear curvature objective function	3×10^4	6×10^2	2×10^1
Quadratic curvature objective function	5×10^3	2×10^2	4×10^{-2}
Penalty curvature objective function	1×10^4	1×10^3	5.0

TABLE 1. Objective function weightings for the four optimization runs. Here w_{Bn} , w_L and w_κ are the field error, length and curvature objective function weightings, respectively.

This forces all local minima to have coil geometries with curvature less than or equal to the input maximum curvature. To obtain coils using $f_{\kappa,3}$ for this letter, the initial value of α is set to 0 and the optimization is run until a local minimum is found. Coils from the initial optimization are then input back into FOCUS with α increased to 0.1 and the optimization is again run until a local minimum is found. This process is repeated with α increasing after each optimization. Coils with a constrained maximum curvature are found after α is increased to 4.

3. Applications

FOCUS is used to solve for coils, targeting an idealized boundary similar to that of the HSX that has been smoothed to remove errors produced by the actual HSX coils. We refer to this smoothed boundary as the HSX boundary. HSX is a four-field-period device with 48 coils. Because of symmetries, there are only six unique coil shapes. The normal field boundary condition is set to zero, $T_{Bn} = 0$, representing a zero beta equilibrium with no fixed external magnets. We initialize the optimizations with filamentary coils that are roughly equal to the centroids of the HSX finite build coil cross-sections. For the remainder of this letter, the term ‘HSX coils’ refers to these filamentary coils.

We present four optimizations to compare the different objective functions: one with no curvature optimization and one for each of the three curvature objective functions. Coils in these optimizations have equal currents of 151 kA and currents are held fixed during optimization. We do not change each coil’s target length between optimizations, we use the nonlinear conjugate gradient algorithm and strong Wolfe conditions (Nocedal & Wright 2006) and we set Fourier series truncation to be $N_F = 16$ for the coils in each optimization. Target coil lengths are chosen to be slightly larger than the HSX coil lengths. The four optimizations are run until local minima are found.

Choice of objective function weightings will significantly change the solution space and resulting coil solutions. Objective function weightings are given in table 1. The value of coil length objective function weightings is chosen such that the difference between the resulting coil lengths and target coil lengths is of the order of centimetres for each of the four optimization runs. Weighting for the linear curvature objective function was adjusted until the field error objective function value f_{Bn} for the resulting coils was close to the f_{Bn} value for the HSX coils. Weighting for the quadratic curvature objective function is set in the same way. It should be noted that the penalty function variable α is more important than the penalty curvature objective function weighting.

To evaluate each optimization result, we compare against the HSX coils. The HSX coils and the four FOCUS-optimized coil sets can be seen in figure 3. A comparison of the coil curvatures can be seen in figure 4. Important metrics to consider for each coil set are the field error objective function, maximum field error, maximum curvature

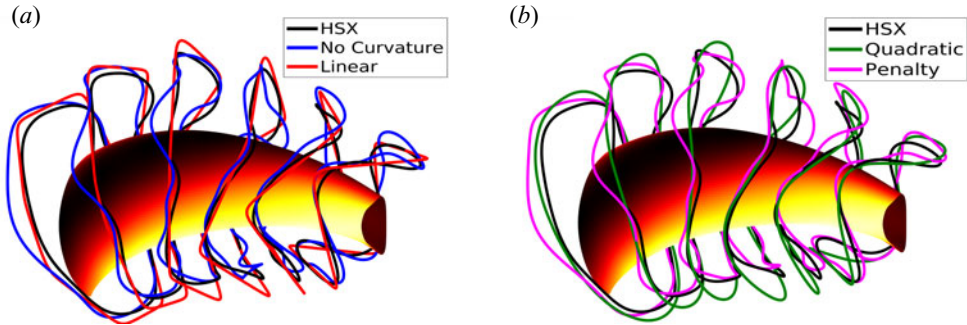


FIGURE 3. Coils optimized for the HSX boundary are shown. HSX coils that are not optimized are plotted in black and are included in both panels. (a) FOCUS coils with no curvature optimization are plotted in blue and FOCUS coils optimized with linear curvature are plotted in red. (b) FOCUS coils optimized with quadratic curvature are plotted in green and FOCUS coils optimized with our penalty curvature objective function are plotted in magenta. Parameter $|\mathbf{B}_v|$ from coils with no curvature optimization is plotted on a half-period of the HSX boundary in (a) and $|\mathbf{B}_v|$ from coils optimized with our penalty curvature objective function is plotted on a half-period of the HSX boundary in (b). High $|\mathbf{B}_v|$ is plotted in black and low $|\mathbf{B}_v|$ is plotted in yellow.

and minimum coil-to-coil separation. These values are given in table 2. For the penalty curvature objective function, the value of κ_0 is set equal to the HSX coil maximum curvature and we can see that our penalty curvature objective function has been used to successfully constrain the coil curvature.

4. Discussion

4.1. No curvature optimization

If $N_F = 16$ and coil curvature is not optimized for, it is evident that FOCUS does not solve for reasonable coils due to high coil curvatures, with radius of curvature in places reduced below 1 cm. One technique to reduce coil curvatures is to decrease the value of N_F . Even for low values of N_F , coil curvature can still become too high. Using a low value of N_F restricts coil shaping and leads to higher field errors.

4.2. Linear curvature optimization

Including the linear curvature objective function produces coils with low average curvature. Unfortunately, the maximum curvature of these coils is 74.4 m^{-1} which is significantly larger than that of the HSX coils that have a maximum curvature of 12.3 m^{-1} , and both coil sets have roughly equal f_{Bn} values. After close inspection it can be concluded that some regions of the coils have been straightened out while other regions of the coils still have sharp bends. High curvatures are not heavily penalized as long as they exist over short arc lengths.

4.3. Quadratic curvature optimization

The quadratic curvature objective function should penalize regions with high curvatures more than does the linear curvature function. Coils generated with the quadratic curvature function have a maximum curvature of 7.62 m^{-1} which is lower than the HSX coil's maximum curvature of 12.3 m^{-1} . Again both coil sets have roughly equal values of f_{Bn} . It can be seen from figure 1(b) that the field error distribution does not appear to be dominated by coil ripple. A large portion of the error comes from non-ripple error. For

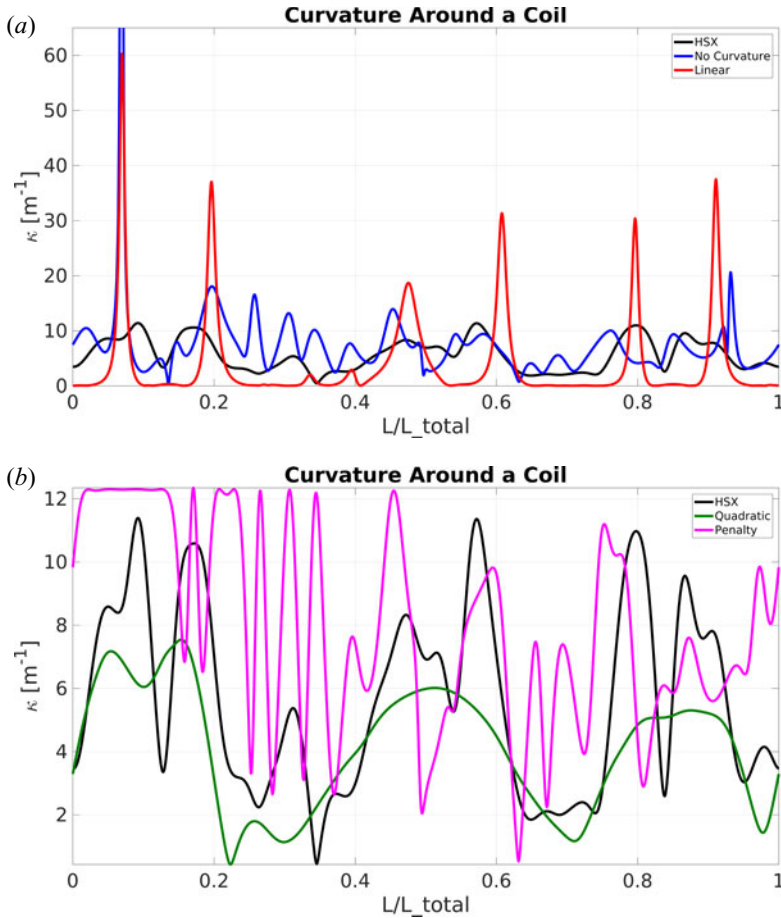


FIGURE 4. Coil curvature is plotted against the length around the fifth coil. HSX coils that are not optimized are plotted in black and are included in both panels. (a) FOCUS coils with no curvature optimization are plotted in blue and FOCUS coils optimized with linear curvature are plotted in red. (b) FOCUS coils optimized with quadratic curvature are plotted in green and FOCUS coils optimized with our penalty curvature objective function are plotted in magenta. The plot in (a) is truncated to a maximum curvature of 65 m^{-1} . The blue spike that is cut off in (a) extends to a maximum value of 246.

roughly equal values of f_{Bn} , it appears that the quadratic curvature objective function outperforms the linear curvature objective function with regards to the maximum curvature. Coils optimized with the linear curvature objective function have a maximum curvature that is ≈ 10 times higher than coils optimized with the quadratic curvature objective function.

4.4. Penalty function curvature optimization

The penalty function is designed to only penalize positions on the coils where the curvature exceeds a prescribed maximum, here set to be the HSX coil curvature maximum, $\kappa_0 = 12.3$. Coils solved for with the penalty function are shown to have a maximum curvature equal to that of the HSX coils. The value of f_{Bn} is more than an order of magnitude lower than that of the HSX coils. Remarkably, the value of f_{Bn} for these coils is

Optimization run	f_{Bn}	$\max \left \frac{\mathbf{B}_v \cdot \mathbf{n}}{ \mathbf{B}_v } \right $	$\max \kappa$	$\min ccsep$
HSX coils	1.49×10^{-4}	2.28×10^{-2}	12.3	9.3
No curvature objective function	6.72×10^{-6}	7.90×10^{-3}	246	8.0
Linear curvature objective function	1.49×10^{-4}	2.28×10^{-2}	74.4	8.3
Quadratic curvature objective function	1.49×10^{-4}	3.37×10^{-2}	7.62	8.2
Penalty curvature objective function	9.92×10^{-6}	1.02×10^{-2}	12.3	8.5

TABLE 2. Important coil metrics for the initial HSX coils and four optimization runs are given. The normalized field error objective function is given as f_{Bn} . Maximum normalized field error is given as $\max |(\mathbf{B}_v \cdot \mathbf{n})/|\mathbf{B}_v||$. Maximum curvature is given as $\max \kappa$ (m^{-1}). Minimum coil-to-coil separation is given as $\min ccsep$ (cm).

close to the value of f_{Bn} from the FOCUS coils that are not optimized for curvature. Since we do not penalize locations of coils that have curvatures below our specified maximum curvature, the coils have more flexibility to accurately reconstruct the plasma boundary. It can be seen from [figure 1\(a\)](#) that the field error distribution appears to be dominated by coil ripple and the normalized ripple error is roughly 0.25%. This means that the normalized non-ripple error is even lower than 0.25% and because of this we have higher confidence that the general shape of the ‘plasma boundary’ is reconstructed well.

5. Conclusions

Without a curvature objective function, optimizations in FOCUS can quickly solve for coils that have unreasonably high curvatures. There are techniques that can be used to limit coil curvatures, but typically at the cost of increasing field errors. Because of this, multiple curvature objective functions have been implemented in FOCUS. We see that a quadratic curvature objective function seems to outperform a linear curvature objective function with respect to maximum coil curvature, but does not outperform with respect to maximum field error. Neither result is satisfactory. This motivated a constraint on maximum curvature as opposed to optimizing for curvature at every location of the coil. To apply this constraint we used penalty functions. The resulting coils with the penalty functions outperformed both the linear and quadratic objective functions with regard to matching the target boundary while still remaining below the maximum curvature threshold.

If coil buildability is our primary concern, then a coil’s maximum curvature seems to be a more apt metric than the average coil curvature or average squared curvature. We see that constraining a coil’s maximum curvature with a penalty function can solve for coils with a very low value of the field error objective function, f_{Bn} . A hyperbolic cosine penalty function was chosen for the optimization mainly due to its exponential penalty and its differentiability.

The penalty objective function achieves the desired behaviour of only penalizing locations of the coils where curvatures exceed a maximum curvature. However, we can see in [figure 3\(b\)](#) that there are small wiggles in the coils produced with the penalty objective function. These deviations may decrease field errors by a small amount, but more importantly, they may also significantly increase ‘coil complexity’. Removing these additional wiggles is a topic that can be addressed with additional curvature objective functions that specifically target coil complexity. This is left for a future publication.

In this letter we demonstrate how to use penalty functions to constrain coil curvature. We intend to also use penalty functions to constrain the minimum coil-to-coil separation. Many more useful applications of penalty functions exist outside of coil optimization. For instance, when optimizing for stellarator equilibria, it is often desirable to constrain rotational transform profiles so as to not cross low-order rational values. Penalty functions are easy to implement and do not require additional optimization algorithms for their implementation.

Acknowledgements

This work was supported by the US Department of Energy (T.G.K., contract number DE-FG02-93ER54222; C.Z., contract number DE-AC02-09CH11466; A.B., contract number DE-FG02-93ER54222; D.T.A., contract number DE-FG02-93ER54222; L.S., contract number DE-FG02-93ER54222) and the University of Wisconsin-Madison (T.G.K., grant number 135AAD3116; A.B., grant number 135AAD3116).

Editor William Dorland thanks the referees for their advice in evaluating this article.

Declaration of interests

The authors report no conflict of interest.

Data availability statement

The data that support the findings of this study are openly available in the HSX documents repository at <https://hsx.wisc.edu/documents/constrained-stellarator-coil-curvature-optimization-with-focus/>.

Appendix A. Objective function derivatives

First derivatives of the three curvature objective functions are now given. Here X_i is an arbitrary Fourier mode amplitude, $X_i \in \{X_{c,0,i}, \dots, X_{c,N_F,i}, X_{s,1,i}, \dots, X_{s,N_F,i}\}$, that defines the x position of the i th coil. Derivatives for Y_i or Z_i Fourier mode amplitudes take a similar form. The equation for $\partial L_i / \partial X_i$ is given in (15) in Zhu *et al.* (2017). The derivatives are

$$\begin{aligned} \frac{\partial f_{\kappa,1}}{\partial X_i} = & -\frac{1}{N_c} \frac{1}{L_i^2} \frac{\partial L_i}{\partial X_i} \int_0^{2\pi} \kappa_i |\mathbf{r}_i'| dt + \frac{1}{N_c} \frac{1}{L_i} \int_0^{2\pi} \frac{\partial \kappa_i}{\partial X_i} |\mathbf{r}_i'| dt \\ & + \frac{1}{N_c} \frac{1}{L_i} \int_0^{2\pi} \kappa_i \frac{\partial |\mathbf{r}_i'|}{\partial X_i} dt, \end{aligned} \quad (\text{A } 1)$$

$$\begin{aligned} \frac{\partial f_{\kappa,2}}{\partial X_i} = & -\frac{1}{N_c} \frac{1}{L_i^2} \frac{\partial L_i}{\partial X_i} \int_0^{2\pi} \kappa_i^2 |\mathbf{r}_i'| dt + \frac{1}{N_c} \frac{1}{L_i} \int_0^{2\pi} 2\kappa_i \frac{\partial \kappa_i}{\partial X_i} |\mathbf{r}_i'| dt \\ & + \frac{1}{N_c} \frac{1}{L_i} \int_0^{2\pi} \kappa_i^2 \frac{\partial |\mathbf{r}_i'|}{\partial X_i} dt, \end{aligned} \quad (\text{A } 2)$$

$$\begin{aligned} \frac{\partial f_{\kappa,3}}{\partial X_i} = & -\frac{1}{N_c} \frac{1}{L_i^2} \frac{\partial L_i}{\partial X_i} \int_0^{2\pi} H_{\kappa_0}(\kappa_i) (\cosh(\alpha(\kappa_i - \kappa_0)) - 1)^2 |\mathbf{r}_i'| dt \\ & + \frac{1}{N_c} \frac{1}{L_i} \int_0^{2\pi} H_{\kappa_0}(\kappa_i) 2\alpha (\cosh(\alpha(\kappa_i - \kappa_0)) - 1) \sinh(\alpha(\kappa_i - \kappa_0)) \frac{\partial \kappa_i}{\partial X_i} |\mathbf{r}_i'| dt \\ & + \frac{1}{N_c} \frac{1}{L_i} \int_0^{2\pi} H_{\kappa_0}(\kappa_i) (\cosh(\alpha(\kappa_i - \kappa_0)) - 1)^2 \frac{\partial |\mathbf{r}_i'|}{\partial X_i} dt. \end{aligned} \quad (\text{A } 3)$$

The i coil indices are dropped for convenience. Now we give the derivative of curvature with respect to an arbitrary x Fourier mode amplitude:

$$\frac{\partial \kappa}{\partial X} = \frac{(\mathbf{r}' \times \mathbf{r}'') \cdot \left(\frac{\partial x'}{\partial X} \hat{\mathbf{x}} \times \mathbf{r}'' + \frac{\partial x''}{\partial X} \mathbf{r}' \times \hat{\mathbf{x}} \right)}{|\mathbf{r}' \times \mathbf{r}''| |\mathbf{r}'|^3} - \frac{3 |\mathbf{r}' \times \mathbf{r}''| x' \frac{\partial x'}{\partial X}}{|\mathbf{r}'|^5}. \quad (\text{A } 4)$$

REFERENCES

DREVLAK, M. 1998 Automated optimization of stellarator coils. *Fusion Technol.* **33** (2), 106–117.

DREVLAK, M. 1999 Optimization of heterogenous magnet systems. In *Proceedings of the 12th International Stellarator Workshop*.

GATES, D. A., ANDERSON, D. VID, ANDERSON, S., ZARNSTORFF, M., SPONG, D. A., WEITZNER, H., NEILSON, G. H., RUZIC, D., ANDRUCZYK, D., HARRIS, J. H., *et al.* 2018 Stellarator research opportunities: a report of the national stellarator coordinating committee. *J. Fusion Energy* **37** (1), 51–94.

LANDREMAN, M. 2017 An improved current potential method for fast computation of stellarator coil shapes. *Nucl. Fusion* **57** (4), 046003.

LANDREMAN, M. & BOOZER, A. H. 2016 Efficient magnetic fields for supporting toroidal plasmas. *Phys. Plasmas* **23** (3), 032506.

MERKEL, P. 1987 Solution of stellarator boundary value problems with external currents. *Nucl. Fusion* **27** (5), 867.

NEILSON, G. H., GRUBER, C. O., HARRIS, J. H., REJ, D. J., SIMMONS, R. T. & STRYKOWSKY, R. L. 2010 Lessons learned in risk management on NCSX. *IEEE Trans. Plasma Sci.* **38** (3), 320–327.

NOCEDAL, J. & WRIGHT, S. 2006 *Numerical Optimization*. Springer Science & Business Media.

RIBBE, K. 2009 Experiences from design and production of Wendelstein 7-X magnets. *Fusion Engng Des.* **84** (7–11), 1619–1622.

SINGH, L., KRUGER, T. G., BADER, A., ZHU, C., HUDSON, S. R. & ANDERSON, D. T. 2020 Optimization of finite-build stellarator coils. *J. Plasma Phys.* **86** (4).

STRICKLER, D. J., BERRY, L. A. & HIRSHMAN, S. P. 2002 Designing coils for compact stellarators. *Fusion Sci. Technol.* **41** (2), 107–115.

ZHU, C., HUDSON, S. R., SONG, Y. & WAN, Y. 2017 New method to design stellarator coils without the winding surface. *Nucl. Fusion* **58** (1), 016008.

ZHU, C., HUDSON, S. R., SONG, Y. & WAN, Y. 2018 Designing stellarator coils by a modified Newton method using FOCUS. *Plasma Phys. Control. Fusion* **60** (6), 065008.