

Use of a genetic algorithm for compact stellarator coil design

W.H. Miner, Jr., P.M. Valanju
Fusion Research Center,
The University of Texas at Austin,
Austin, Texas

S.P. Hirshman
Oak Ridge National Laboratory,
Oak Ridge, Tennessee

A. Brooks, N. Pomphrey
Princeton Plasma Physics Laboratory,
Princeton University,
Princeton, New Jersey

United States of America

Abstract. A new global optimization technique for designing stellarator coils has been developed and applied to the design of coils for the National Compact Stellarator Experiment. Using this technique coil sets were found with fewer coils and lower current densities than those obtained with traditional methods. A new coil design procedure which uses a genetic algorithm as the core optimization method is described and the resulting advanced coil designs presented.

1. Introduction

The National Compact Stellarator Experiment (NCSX) has been proposed for investigation of the physics of high beta, low aspect ratio, quasi-axisymmetric (QA) stellarator plasmas. An NCSX plasma can be regarded as a specific helical perturbation of an axisymmetric tokamak plasma. The required perturbation is such that on each flux surface the magnetic field strength, when expressed as a function of magnetic (Boozer) co-ordinates, is independent of toroidal angle. QA stellarators offer the combined advantages of tokamak neoclassical confinement levels and stellarator stability.

The first stage of NCSX design requires identification of a baseline plasma configuration with attractive physics properties. A three field period QA equilibrium with an aspect ratio of 3.5 that is stable to kink and ballooning modes at $\beta = 4\%$ has been obtained [1] using a configuration optimization procedure developed by Nührenberg and Zille [2]. Fourier coefficients of the plasma boundary shape were varied to minimize a target function which expresses the desired physics properties of the NCSX plasma. The target function includes values of

the rotational transform at the magnetic axis and the plasma edge (global shear), a measure of the degree of quasi-symmetry (goodness of transport), measures of kink and ballooning stability, and boundary limits describing geometric constraints such as aspect ratio, maximum or minimum radius and maximum plasma height. The VMEC [3] MHD equilibrium code computes a new equilibrium for each boundary modification, from which the optimizer target function can be evaluated. When one or more attractive equilibria have been identified the process of designing coils to support the equilibria can begin.

In this article we will focus on a coil design for a baseline equilibrium denoted C82. The last closed flux surface of C82 is shown in Fig. 1. On this surface the coil set must impose the condition that the normal component of the total magnetic field must vanish,

$$\mathbf{B} \cdot \hat{\mathbf{n}} = 0 \quad (1)$$

where

$$\mathbf{B} = \mathbf{B}_{pl} + \mathbf{B}_{ext}. \quad (2)$$

Here $\hat{\mathbf{n}} = \nabla s / |\nabla s|$ is the unit normal to the plasma bounding surface, $s(x, y, z)$, \mathbf{B}_{pl} is the magnetic field




Figure 1. A three dimensional rendering of the C82 plasma boundary.

due to currents which flow in the plasma and \mathbf{B}_{ext} is the magnetic field due to all external field coils. There is a possibility for building the NCSX device on the existing PBX platform and of making use of the PBX toroidal and poloidal magnetic field coils. It is convenient, therefore, to write \mathbf{B}_{ext} as the sum

$$\mathbf{B}_{ext} = \mathbf{B}_{tok} + \mathbf{B}_{stell} \tag{3}$$

where \mathbf{B}_{tok} is the (known) magnetic field from any given tokamak coils, and \mathbf{B}_{stell} is the magnetic field from the stellarator coils whose design is being sought.

Following the NESCOIL [4] procedure developed by Merkel, \mathbf{B}_{stell} can be produced by a surface current \mathbf{j} on a current winding surface (CWS) that encloses the plasma. The surface current is divergence free and can be written in terms of the surface gradient ∇_s of a current potential $\Phi(u, v)$,

$$\mathbf{j} = \nabla s' \times \nabla_s \Phi(u, v). \tag{4}$$

Here $s'(x, y, z)$ defines the shape of the CWS and u, v are poloidal, toroidal angles on that surface. Once the potential is determined [4], a set of coils can be defined by selecting N_c appropriate contours of Φ and interpreting each contour as a filamentary coil carrying an amount of current that is proportional to the change in potential midway between each chosen contour and its two chosen neighbours. In the limit as $N_c \rightarrow \infty$ the discrete coil system reproduces exactly the magnetic field of the current sheet. For a practical coil system, however, we would like to choose an efficient coil set with the following minimum set of properties:




Figure 2. A contour plot of the surface current potential for C82. Solid lines represent contours of positive potential, short dashed lines represent contours of negative potential and long dashed lines represent contours of zero potential.




Figure 3. Cross-sections of the plasma boundary and the coil winding surface in the two toroidal planes $N\phi = 0$ and π .

- (a) The number of coils should be small.
- (b) The reconstruction errors (measured by how well Eq. (1) is satisfied) should be small.
- (c) The maximum coil current should be small to minimize resistive dissipation.

An optimization algorithm for choosing the ‘best’ set of N_c contours is the primary topic of this article.

Figure 2 shows a contour plot of the NESCOIL current potential for C82 on a CWS (Fig. 3) that

is conformal with the plasma boundary, with each point of the CWS a normal distance of 18 cm from the plasma. In NCSX this normal separation is determined by engineering constraints: it is the smallest coil–plasma separation that allows room for plasma facing components, vacuum vessel, etc., in an $R = 1.5$ m device. As the separation distance is increased, both the normal magnetic field error at the plasma boundary and the current sheet density on the CWS increase. The simplest algorithm for cutting coils from the current sheet is to choose N_c contours equally spaced in Φ , with equal currents in each of the coils. This, in fact, was the algorithm used in the coil designs for Helias and ATF [4]. However, the NCSX plasma has a smaller aspect ratio and is more highly shaped (owing to the requirements of quasi-symmetry and high beta stability), and it becomes extremely beneficial to optimize the coil selection procedure: the number of stellarator coils and the maximum current required by the coils can be significantly decreased.

Initially, an optimization code was written which selected coil contours on the basis of minimizing a weighted sum of the squares of the net normal component of the magnetic field at the plasma boundary (which, ideally, should vanish) and the maximum current in the selected coils. A Levenberg–Marquardt algorithm [5] was used as the minimization procedure. Modest improvements in coil design were obtained when compared with the simple equipotential spacing algorithm (smaller B_{nor} fitting errors and 10–20% smaller maximum coil currents for a fixed number of coils). However, it became clear that the final Levenberg–Marquardt solutions were trapped in local minima of the minimizing target function and that considerable gains might be achieved if a global minimization algorithm were applied.

In the coil selection problem we are confronted with the daunting combinatorial task of selecting a relatively small number of potential contours (coils) from a large pool of contours (the full current sheet) and requiring that they reconstruct the plasma boundary with sufficient accuracy to reproduce the desired physics of the original equilibrium. A specific example will demonstrate the enormity of the problem. In Fig. 2 there are 60 equally spaced contours of the C82 current sheet. These produce 98 individual coils (some levels are multivalued, as can be seen from the saddle region of the figure). If we seek a coil design with only 10 coils per field period, then

all combinations of these 98 coils taken 10 at a time — which leads to

$$\frac{98!}{10!(98-10)!} = 1.4 \times 10^{13} \text{ combinations!}$$

— would have to be examined to guarantee finding a global minimum of the target function (minimum B_{nor} , together with any current density minimization).

Thus, the combinatorial optimization problem confronting us consists of a ‘state space’, a target ‘cost function’ to be minimized and a discrete but very large set of possible solutions. From previous calculations, we also know that the solution space exhibits many local minima. One approach to this problem is to apply a genetic algorithm (GA) [6], which is a non-gradient based optimization method that has been successfully applied to several problems of this type [7]. The GA works with a population of ‘individuals’, each of which represents a possible solution to the optimization problem. Each individual is assigned a ‘fitness’ according to the value of the cost function. The fittest individuals, those with the lowest value of the cost function, are allowed to reproduce by cross-breeding, thereby producing a new generation of individuals (population of new possible solutions) that contains a high proportion of the best characteristics of the previous generation. In this way, over generations, the good characteristics (low cost) are spread throughout the population and the most promising areas of the search space can be explored in a computationally efficient manner.

2. The genetic algorithm

2.1. Overview

An overview of the steps in the GA is shown in Fig. 4. For further information and different applications the interested reader is referred to Refs [8–12].

To apply the GA to the problem of cutting discrete coils for a stellarator, we proceed as follows. First the current potential $\Phi(u, v)$ defining a current sheet on a specified current winding surface is discretized into a large number of equally spaced contours (e.g. Fig. 2). These contours represent the pool of individuals from which the final, much smaller, subset of coils will be chosen. The GA will be used to decide which coils are to be retained, on the basis of their evaluated fitness (matching B_{err}^{avg} and current density targets). Let N_Φ denote the total number of coil contours. In addition, let N_c denote the number of coils desired in the final coil design.

	binary string
jump mutation before	010010010001000
jump mutation after	010000010001010
creep mutation before	010100001000100
creep mutation after	010100010000100

Figure 6. Examples of the two mutation processes.

The completion of the steps outlined in Fig. 4 constitutes a generation. The procedure is repeated, beginning with a new population, until a satisfactory solution is obtained or a stated maximum number of generations has been examined.

2.2. Fixing the number of coils in a set

In the usual GA, the range of the variable represented by the binary string is set such that the minimum value is given by the string consisting of all zeros (all bits are ‘off’), and the maximum value is given by the string consisting of all ones (all bits are ‘on’). Therefore, in the usual GA any number of bits may be ‘on’ in any chromosome. In the case of the coil selection process, a fixed number of coils N_c is sought. Therefore, an additional step has to be taken in order to preserve this number.

After the GA has produced a new individual, it is checked to determine its length, i.e. the number of bits that are ‘on’. If there are more than N_c bits ‘on’ the chromosome should either be discarded or else be modified. Since discarding the individual would then require an additional mating, and we would not be guaranteed that the new chromosome would be too long, we chose to modify the chromosome. This was done using the following procedure. If a bit is ‘off’, the next bits are scanned until an ‘on’ bit is found. A random number $0 \leq P_{change} \leq 1$ is generated. If $P_{change} \leq 0.5$ the bit is turned ‘off’ and if $P_{change} > 0.5$ it remains unchanged. Then, starting at the next bit the procedure is repeated until the number of bits that are ‘on’ equals N_c . Note that if a chromosome has less than N_c bits ‘on’ it is allowed in the population because coil sets with fewer coils are desirable.

3. Computational details

3.1. The inductance matrix

Previous methods for finding stellarator coils have used a Fourier representation for the shape of

filamentary coils on the CWS and a Fourier representation for the shape of the CWS itself [14]. The Fourier coefficients of both the filaments and the CWS are allowed to vary in an optimization loop which minimizes a combination of engineering constraints. Such techniques are very flexible, but the calculation of the magnetic field at the plasma boundary is computationally expensive. Each time the coils are moved the magnetic field on the plasma boundary must be recomputed using the Biot–Savart law.

For the GA, the coil pool from which the evolving population is constructed is fixed. An inductance matrix \mathbf{G} , of dimension $N_{uv} \times N_\Phi$ relates currents in each of the N_Φ coils in the pool to the normal magnetic field at N_{uv} points on the plasma boundary,

$$\mathbf{G} \cdot \mathbf{I} = \hat{\mathbf{n}} \cdot \mathbf{B}_{stell}. \quad (5)$$

The full \mathbf{G} matrix can be calculated once outside of the optimization loop, then reused within the loop when calculating the magnetic field on the plasma boundary due to each member of the population in each generation. This simply requires compacting the full inductance matrix into an $N_{uv} \times N_c$ reduced matrix \mathbf{G}_c obtained by eliminating those columns of \mathbf{G} which correspond to coils that are not energized. The magnetic field is then evaluated by the simple matrix multiplication

$$\mathbf{G}_c \cdot \mathbf{I}_c = \hat{\mathbf{n}} \cdot \mathbf{B}_{stell}. \quad (6)$$

3.2. Calculation of coil currents

After the GA has chosen to energize a particular coil subset, the current in each of these N_c coils must be evaluated. The choice is that vector \mathbf{I}_c which minimizes

$$\chi^2 = |\mathbf{G}_c \cdot \mathbf{I}_c - \hat{\mathbf{n}} \cdot (\mathbf{B}_{tok} + \mathbf{B}_{pl})|^2. \quad (7)$$

The solution to this subsidiary minimization problem is obtained by singular value decomposition (SVD) [5], giving

$$\mathbf{I}_c = \mathbf{V}[\text{diag}(1/w_j)][\mathbf{U}^T \cdot \hat{\mathbf{n}} \cdot (\mathbf{B}_{tok} + \mathbf{B}_{pl})] \quad (8)$$

where

$$\mathbf{G}_c = \mathbf{U}\mathbf{w}\mathbf{V}^T \quad (9)$$

is the SVD of the compressed inductance matrix. \mathbf{U} and \mathbf{V} are $N_{uv} \times N_c$ and $N_c \times N_c$ unitary matrices (eigenfunctions, respectively, of $\mathbf{G}\mathbf{G}^T$ and $\mathbf{G}^T\mathbf{G}$), and \mathbf{w} is an $N_c \times N_c$ diagonal matrix of ‘singular values’ w_j (the square roots of eigenvalues of $\mathbf{G}\mathbf{G}^T$ and $\mathbf{G}^T\mathbf{G}$).

Table 1. NESCOIL current sheet solution and results using 13 coils per half-period, chosen equally spaced in Φ

Source	Solution	B_{err}^{avg} (%)	B_{err}^{max} (%)	J^{max} (kA/cm ²)
NESCOIL	Current sheet	0.22	2.0	
Equi- $\Delta\Phi$	13 coils per half-period	0.95	7.0	14.7

4. Results

4.1. Coil studies

In general, the current potential on a prescribed CWS enclosing a chosen plasma can be written as the sum of a secular and a periodic contribution [4]:

$$\Phi(u, v) = c_t u + c_p v + \sum_{m=0}^M \sum_{n=-N}^N \Phi_{mn} \sin 2\pi(mu + nv). \quad (10)$$

Non-zero secular coefficients c_t and c_p correspond to coil topologies with net poloidal and toroidal currents (e.g., stellarator modular coils or wavy PF coils). A saddle coil design with $c_t = c_p = 0$ is being considered for NCSX. For this design the toroidal flux requirements of the plasma are provided by external tokamak field coils. The saddle coils provide the required stellarator magnetic field for producing the rotational transform.

We seek a saddle coil design to support an equilibrium configuration named C82 (Fig. 1). The major plasma parameters for C82 are $\langle R \rangle = 1.45$ m, $\langle a \rangle = 0.42$ m, $\beta = 4\%$, $I_p = 200$ kA and $B_T = 1.2$ -2T. A current sheet solution was determined on a CWS conformal with the plasma boundary and separated from it by a normal distance of 18 cm. $N_\Phi = 60$ contours of the current potential are shown in Fig. 2. These contours form the coil pool for the GA.

4.2. Benchmark for comparison of GA results

In order to evaluate the effectiveness of the GA, it is useful to establish a benchmark with which to compare the GA results. We begin by noting two criteria that are needed for an efficient coil design. The first is an engineering constraint: in order to operate the coils reliably and cool them effectively, the coil current density must be kept below 15 kA/cm². This puts a constraint on the minimum allowable separation between coils. The second criterion is a physics constraint: the field produced by the coils must provide a satisfactory match to the field from




Figure 7. Average error obtained by the GA as a function of the number of generations for 12 coils per field period.

the target equilibrium in order for that target to be reconstructed. The average error is

$$B_{err}^{avg} \equiv \frac{1}{N} \sum_{i=1}^N |\mathbf{B}_i \cdot \hat{\mathbf{n}}| / |\mathbf{B}_i| \quad (11)$$

where N is the total number of points at which the field is calculated on the plasma boundary. In practice, it has been found that the average fitting error B_{err}^{avg} must be $\lesssim 1\%$ for an NCSX plasma.

The benchmark was obtained by taking a fixed number of current potential contours, equally spaced in Φ , with equal currents in all coils, and finding the minimum number of coils required to obtain $B_{err}^{avg} < 1\%$. In order to satisfy this limit, 26 coils per field period were required, carrying a maximum current of 14.7 kA/cm². This result is tabulated in Table 1, where, for comparison, we also show the fitting error obtained from the continuous NESCOIL current sheet solution (a lower bound for any discrete coil calculation).

Our initial GA calculation selected 24 coils per period and ran for 3000 generations. The results are




Figure 8. The two helical coils obtained from the zero value of the current potential.

shown in Fig. 7. Although 3000 generations were computed, no better solution was found after generation 1600, the last point on the curve. The GA found a solution with average $B_{err}^{avg} = 0.33\%$ and maximum $B_{err}^{max} = 2.55\%$, an improvement over the similar equally spaced coil set of nearly a factor of three. There are two important points to be noted in this figure. First is the rapid rate of convergence (the number of generations is shown on a logarithmic scale). Second is the accuracy attained. As previously mentioned, good reconstructions have been obtained with errors $\lesssim 1\%$, and the GA has found solutions well below that limit.

The goal of this work is to find the minimum number of coils with current densities below the engineering limit that can still reproduce the magnetic field at the plasma boundary accurately enough that the target plasma can be reconstructed. The results shown in Fig. 7 suggest that we have sufficient margin in B_{err}^{avg} to decrease the number of coils and still meet the 1% error criterion.

4.3. Decreasing the number of coils

Depending on whether the original coil pool is obtained by contouring the current potential with an odd or an even number of contours, a helical coil associated with the contour value $\Phi = 0$ may or may not be present in the coil pool. The $\Phi = 0$ contour is shown in Fig. 8 for the three periods of the CWS. Each helical coil wraps around the plasma twice in the toroidal direction (V) and three times in the poloidal direction (U) before joining back on itself. The helical coil is clearly a special topology among the saddle coils. Since its effect on reconstructing the target plasma is not known a priori, it was deemed important to include — at least initially — the zero potential coil in the available pool.

Thus, beginning with the NESCOIL current potential we choose two sets of contours to provide the candidate coils, namely those obtained by choosing 60 and 61 equally spaced contour levels. In each case, because of the coil degeneracies previously mentioned in the Introduction, a total of 98 individual coils are in each pool. In the case of 60 equally spaced contour levels, this leads to 49 independent coils and their 49 stellarator symmetric images. In the case of 61 equally spaced contour levels, there are 50 independent coils. There are 48 saddle coils and their stellarator symmetric images and the 2 helical coils described earlier.

We begin by targeting only B_{err}^{avg} in the cost function and ignore the engineering constraint on the current density. The GA was asked to select 14 coils per period from the set of possible candidates that minimize B_{err}^{avg} . The GA ran for 3000 generations and found a set of coils that reduced B_{err}^{avg} to $\approx 0.5\%$, a little more than twice the theoretical minimum NESCOIL current sheet solution error of 0.22%. The number of coils was then reduced by one and the process repeated until the maximum acceptable error of 1% was exceeded. The results are shown in Fig. 9. For between 8 and 14 coils per period the error decreases almost linearly. Only in the case of 6 coils per period does the error exceed 1%. None of the final coil solutions contained a helical coil.

4.4. Targeting a combination of B_{err}^{avg} and J_c^{max}

Almost all of the coil solutions shown in Fig. 9 have maximum coil current densities J_c^{max} that exceed 15 kA/cm^2 . To obtain coil solutions that satisfy both the physics reconstruction and the engineering coil current density criteria it is necessary to target both B_{err}^{avg} and J_c^{max} in the GA cost function.

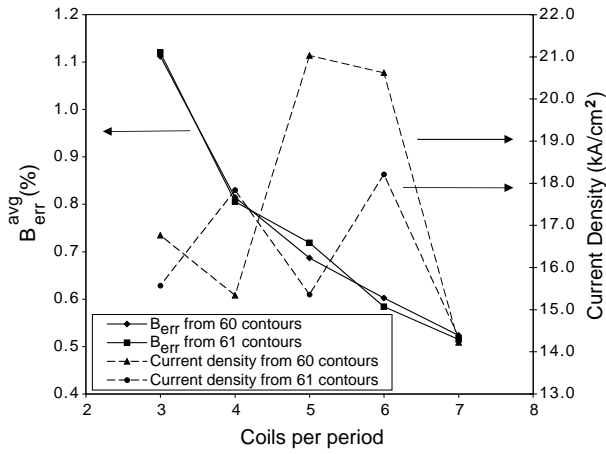


Figure 9. Average error and current density as a function of the number of coils per half period.




Figure 10. Variation of minimum coil separation with the curvature of the current carrying surface.

In this article, we use a crude estimate for the coil current density. It approximates the finite extent (‘build’) of the coils, as well as the coil supports (‘ligaments’). A schematic diagram of this is shown in Fig. 10. Let i and j denote neighbouring coils, carrying coil currents I_i and I_j . The coil centroids lie on the CWS and are separated by a distance d_{sep} . Neighbouring coils are assumed to have the same width w and the same height h . Finally, let $w_{support}$ denote the width of the coil support located between the coils labelled i and j . On the CWS there are as many neighbouring coil pairs as there are individual coil members. An averaged coil current density is calculated for the coil pair ij by averaging the currents over the adjacent coil blocks,

$$J_{ij} = \frac{0.5|I_i + I_j|}{h_{copper}(d_{sep} - w_{support})}. \quad (12)$$

The maximum coil current density is then

$$J_c^{max} = \max_{[ij]} J_{ij}. \quad (13)$$

The coil build and support parameters assumed are $h_{copper} = 7.0$ cm and $w_{support} = 1.0$ cm. The coil

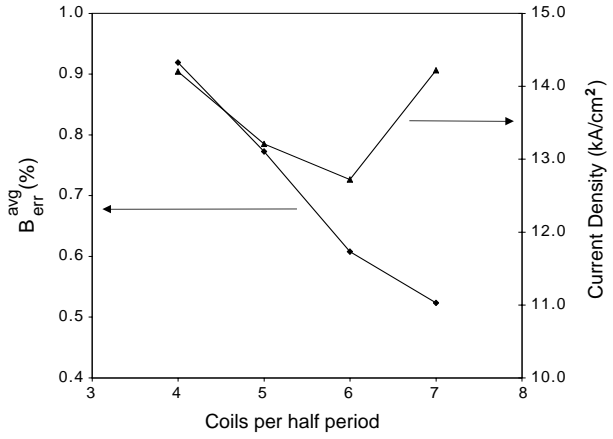


Figure 11. The minimum error obtained where $J_{max} \leq 15$ kA/cm².

Table 2. Comparison of the uniform $\Delta\Phi$ equal current case with the GA results (CPFP stands for coils per field period)

Source	CPFP	(B_{err}^{avg})	(B_{err}^{max})	J_{max}
Equal	26	0.95	7.0	14.7
GA	14	0.52	2.8	14.2
GA	12	0.61	3.8	12.7
GA	10	0.77	5.7	13.2
GA	8	0.92	5.0	14.2

currents and separation distances are calculated by the GA for each candidate coil solution (population member).

The composite target cost function can be written as

$$\mathcal{C} = (1.0 - \alpha)B_{err}^{avg} + \alpha J_c^{max} / J_{norm} \quad (14)$$

where J_{norm} is a scale factor used to normalize the magnitude of the current density so that it will be of the same order as the error. In the previous section we used $\alpha = 0$, and therefore current density played no part in the determination of the error. Now, since we are looking for solutions with $J_c^{max} \leq 15$ kA/cm², we scan α between 0 and 1. The lowest average B errors with current densities below this upper bound are shown in Fig. 11.

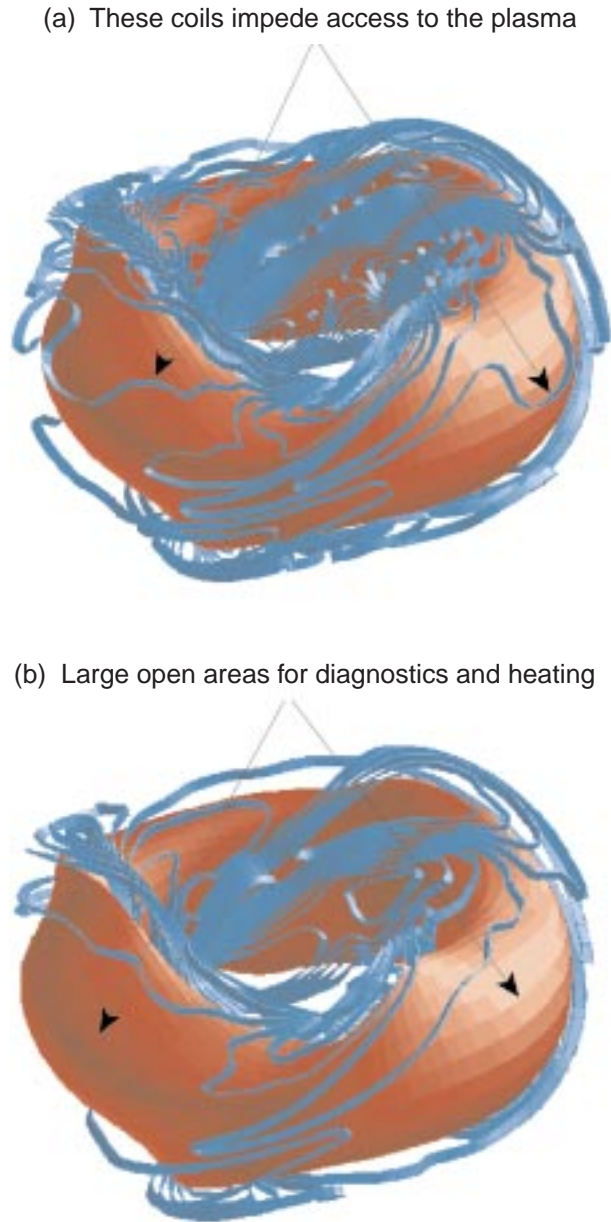


Figure 12. Comparison between the coil set found using equally spaced contours (a) and the coil set found by the GA (b) both having almost equal error and current density.

5. Conclusions

The GA is an extremely useful and flexible tool for finding coils that meet the engineering constraints for a compact stellarator. In this article the GA was used to obtain coil sets with both fewer coils per period and lower current densities than had been obtained using previous coil cutting algorithms [4]. The final

results are compared in Table 2 with the uniform $\Delta\Phi$ case (corresponding to equal currents in all the saddle coil filaments). The coil set obtained using the GA also provides better access to the plasma by reducing the number of coils while still maintaining a constant value of B_{err}^{avg} . Figure 12 shows the coil set found by using the equally spaced contours (26 coils per period) and a coil set found by the GA (8 coils per period). The red surface represents the plasma boundary. In particular, note that the GA was able to eliminate the coil that previously obstructed access at the outboard midplane.

In this article we have demonstrated the usefulness of the GA for determining filamentary saddle coil sets for low aspect ratio stellarators. The next step in the design of a real coil set will be to assess the impact of finite coil thickness on the B_{err}^{avg} matching process. Although this is a straightforward generalization of the procedure described in this article, its implementation and implications are beyond the scope of the present work. In addition, the adequacy of the level of B_{err}^{avg} attained with any set of coils must be determined by comparing the magnetic surfaces reconstructed from the coils with the set of surfaces originally obtained from a stellarator optimization process. This comparison is shown in Fig. 13, which shows the original flux surfaces (solid curves) for C82 and those obtained from the free boundary VMEC code, using the saddle coils shown in Fig. 12. The values of the average and maximum displacements between the original plasma boundary and the boundary obtained from the free boundary equilibrium calculation using the filamentary coils are shown in Table 3. Although the displacement for the GA obtained filaments is larger by a factor of about two, the number of coils in the GA set is smaller by a factor of over three. In addition, the larger error for the GA coil set can be seen in the figure to be localized to the points in the crescent shape and the outboard midplane in the $N\phi = 0$ plane. All other regions of the surface are well fitted.

The ultimate test of the adequacy of any set of coils is to check for the maintenance of the physics properties (transport and stability properties) when a magnetic reconstruction is performed using a free boundary MHD code, such as PIES [15], which does not assume the existence of nested magnetic surfaces. Typically, when this is done it is found necessary to add small correction (or ‘trim’) coils in order to cancel small resonant components of the residual B_{err} spectrum. While the details of this procedure go beyond the scope of the present considerations,

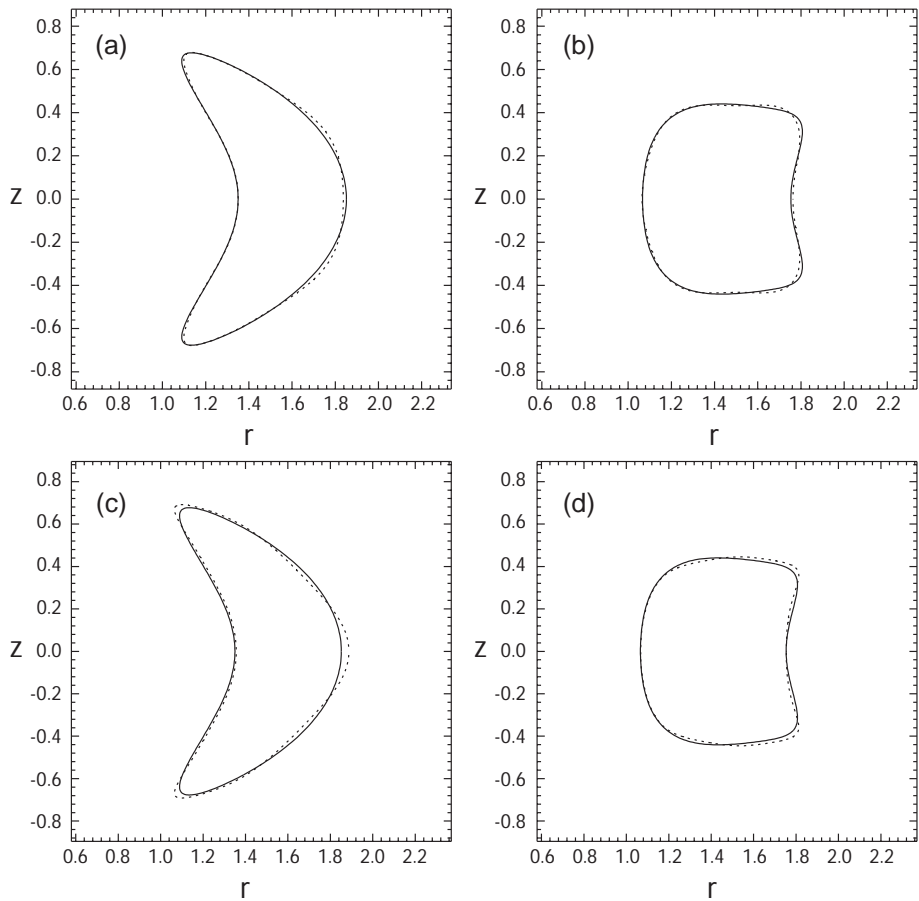


Figure 13. Comparison of the reconstructions using the coil set found using equally spaced contours (a,b) and the coil set found by the GA (c,d), both having almost equal errors and current densities in the two toroidal planes $N\phi = 0$ and π .

Table 3. Comparison of the average displacement and the maximum displacement between the original plasma boundary and the reconstructed plasma boundary for the constant $\Delta\Phi$ constant current case and the eights coils per period GA case.

Source	Average displacement (cm)	Maximum displacement (cm)
Equal	0.47	2.09
GA	1.16	4.37

it is worthwhile mentioning that such a procedure is under development [16].

Acknowledgements

One of the authors (WHM) would like to thank J.C. Wiley for suggesting the GA as a technique

for solving this problem and D.L. Carroll for a copy of his GA code, which he has made available at <http://www.staff.uiuc.edu/~carroll/ga.html>. This work was supported by the US Department of Energy under Grant No. DE-FG03-95ER54296 and Contract Nos. DE-AC020-76-CHO3073 and DE-AC05-00OR22725.

References

- [1] Reiman, A., et al., Plasma Phys. Control. Fusion **41** (1999) B273.
- [2] Nührenberg, J., Zille, R., Phys. Lett. A **129** (1988) 113.
- [3] Hirshman, S.P., van Rij, W.L., Merkel, P., Comput. Phys. Commun. **43** (1986) 143.
- [4] Merkel, P., Nucl. Fusion **27** (1987) 867.
- [5] Press, W.H., et al., Numerical Recipes: The Art of Scientific Computing, Cambridge University Press, Cambridge (1990) 523.
- [6] Holland, J.H., Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence, Ann Arbor, MI (1975).
- [7] Goldberg, D.E., Genetic Algorithms in Search, Optimization, and Machine Learning, Addison-Wesley, Reading, MA (1989).
- [8] Mitchell, M., An Introduction to Genetic Algorithms, MIT Press, Cambridge, MA (1996).
- [9] Haupt, R.L., Haupt, S.E., Practical Genetic Algorithms, Wiley, New York (1998).
- [10] Vose, M.D., The Simple Genetic Algorithm, MIT Press, Cambridge, MA (1999).
- [11] Michalewicz, Z., Genetic Algorithms + Data Structures = Evolution Programs, Springer-Verlag, Berlin (1996).
- [12] Falkenauer, E., Genetic Algorithms and Grouping Problems, Wiley, New York (1998).
- [13] Carroll, D.L., AIAA J. **34** (1996) 338.
- [14] Drevlak, M., Fusion Technol. **33** (1998) 106.
- [15] Reiman, A., Greenside, H., Comput. Phys. Commun. **43** (1986) 157.
- [16] Hudson, S.R., Monticello, D.A., Reiman, A.H., Reduction of Islands in Full-pressure Stellarator Equilibria (in preparation).

(Manuscript received 6 November 2000
Final manuscript accepted 6 April 2001)

E-mail address of W.H. Miner, Jr.:
miner@mail.utexas.edu

Subject classification: B0, Sd