

Exploration of stellarator configuration space with global search methods

H. E. Mynick, N. Pomphrey, and S. Ethier

Plasma Physics Laboratory, Princeton University, P.O. Box 451, Princeton, New Jersey 08543-0451

(Received 6 September 2001; accepted 3 December 2001)

An exploration of stellarator configuration space $\mathbf{z}$ for quasi-axisymmetric stellarator (QAS) designs is discussed, using methods which provide a more global view of that space. To this end, a “differential evolution” search algorithm has been implemented in an existing stellarator optimizer, which is much less prone to become trapped in local, suboptimal minima of the cost function χ than the local search methods used previously. This search algorithm is complemented by mapping studies of χ over $\mathbf{z}$ aimed at gaining insight into the results of the automated searches. It is found that a wide range of the attractive QAS configurations previously found fall into a small number of classes, with each class corresponding to a basin of $\chi(\mathbf{z})$. Maps are developed on which these earlier stellarators can be placed, the relations among them seen, and understanding gained into the physics differences between them. It is also found that, while still large, the region of $\mathbf{z}$ space containing practically realizable QAS configurations is much smaller than earlier supposed. © 2002 American Institute of Physics. [DOI: 10.1063/1.1445756]

I. INTRODUCTION

The search for attractive stellarator designs has been greatly enhanced by the development of optimization codes which search configuration space $\mathbf{z}$ using a specified cost function $\chi(\mathbf{z})$. Such codes have been used extensively in the design of W7-X (Wendelstein-VII-X),¹ HSX (Helically Symmetric experiment),² and more recently, by the NCSX (National Compact Stellarator experiment) group³ in designing a proposed quasi-axisymmetric stellarator (QAS). These optimizers have used search algorithms which are “local,” i.e., which make use of derivatives in the local topography of $\chi(\mathbf{z})$ to decide in which direction to move along a single trajectory in $\mathbf{z}$, such as the steepest-descent and related Levenberg-Marquardt (LM) methods.⁴ The NCSX optimizer (Stellopt) uses the LM algorithm. While efficient in suitable cases, such methods often become trapped in local, suboptimal minima of the cost function, in the large and sometimes highly corrugated $\mathbf{z}$ space. This makes human involvement an essential part of the optimization loop, in which the system $\mathbf{z}$ position or the weights in χ are adjusted to dislodge Stellopt from local minima. Surmounting this difficulty is a first objective of this work, performing more global searches in $\mathbf{z}$ space through implementing a variant (Stellopt-DE) of Stellopt, which uses a “differential evolution” (DE) algorithm.⁵ This is far less prone to become locally trapped, and thus permits a more robust, though more time consuming, exploration of $\mathbf{z}$ space.

The DE algorithm is similar to genetic algorithms (GAs),⁶ but suited to exploration of a continuous space. Unlike local methods, these do not require taking derivatives, but evolve a sequence of generations $g=0, \dots, g_{\max}$, each generation comprised of an ensemble of N_p system points $\{\mathbf{z}^i\}(g), (i=0, \dots, N_p-1)$ distributed over the space to be explored, and with a simple rule for obtaining the $(g+1)^{\text{th}}$ generation from the g^{th} . Because the evolution is of a cloud of system points, rather than of a single point along one

trajectory, the evolution of the DE population can provide a less myopic map of the $\mathbf{z}$ space topography compared with that from traditional local algorithms.

Obtaining such a map of $\mathbf{z}$ space is important, both for understanding and for finding optimal designs. Up to now, the LM optimizer (Stellopt-LM) has obtained promising QAS configurations after a complicated sequence of optimizer runs and human adjustments, and there has been little knowledge of how close the various configurations are to each other, how many different types of good QAS designs there may be, and what their distinguishing features are. In addition to providing an optimization algorithm less prone to the local-well weakness of local methods, a second objective of this work is to enhance insight of this sort, by providing a more global view of the space in which Stellopt has been searching and identifying configurations. We find (cf. Secs. IV and V) that a wide range of the attractive QAS configurations previously studied falls into a small number of classes, with each class corresponding to a well or “basin” of χ . By mapping $\chi(\mathbf{z})$ over useful sections of the full $\mathbf{z}$ space, we obtain maps showing these basins, their relative positions, and gain insight into the physics separating them. Configurations known to have related ancestry tend to fall near one another on the map, as one might expect. The maps also indicate that, while the full $\mathbf{z}$ space is in principal enormous, the region of that space containing configurations which are (in a sense to be defined) practically realizable is much smaller than one might have supposed (though still large). While these maps are not complete or definitive, they begin to put limits on the domain being explored, and to convert *terra incognita* into something with a known geography.

The remainder of this paper is organized as follows. In Sec. II we introduce some needed notation, and briefly discuss features of the operation of Stellopt and ancillary codes which are used for the optimization process. The DE algorithm is introduced in Sec. III, and some of its characteristics

discussed. In Sec. IV the method is employed in studying the QAS region of stellarator configuration space. The search leads to a study of the topography of χ in $\mathbf{z}$ space, and to the development of a simple taxonomy for the QAS configurations on which the NCSX team has focussed. Section V provides some discussion of the results of the foregoing sections.

II. PRELIMINARIES

The cost function $\chi(\mathbf{z})$ is given in Stellopt by

$$\chi^2 = \sum_i \chi_i^2 = \sum_i w_i \hat{\chi}_i^2, \quad (1)$$

where the w_i are the weights of the various contributions $\hat{\chi}_i^2$ to χ^2 . To compute χ^2 , Stellopt first calls the VMEC code⁷ to compute a magnetohydrodynamic (MHD) equilibrium, and then additional codes to compute the different $\hat{\chi}_i^2$, which make use of the equilibrium information. VMEC can compute both fixed and free-boundary equilibria. In the fixed-boundary case, $\mathbf{z}$ specializes to the set $\mathbf{X} = \{X_j = 1, \dots, N_x\} \equiv \{R_{m_1}, Z_{m_1}, R_{m_2}, Z_{m_2}, \dots\}$ of Fourier amplitudes defining the plasma boundary $[R(\theta, \zeta), Z(\theta, \zeta)]$, plus a small number of other quantities such as central pressure p_0 and enclosed toroidal flux Φ_a needed to satisfy certain constraint targets, such as the desired plasma β and toroidal field. Here, $\mathbf{m} \equiv (m, \tilde{n} = n/N_f)$ are poloidal and toroidal mode numbers per period, with N_f equal to the number of field periods, θ and ζ are the poloidal and toroidal azimuths, R is the major radius and Z is the height above the midplane. In the free-boundary case, $\mathbf{z}$ may be given by the set $\mathbf{I} = \{I_j\}$ of currents in a given coil set, plus the few additional quantities just mentioned.

Various transport figures of merit have been used in Stellopt. A simple and standard one often used for studying QA configurations, employed in the present work, is the quasi-symmetry measure $\chi_{Bmn}^2 \equiv \langle \sum_{m,n \neq 0} B_{mn}^2 / B_{00}^2 \rangle_s$, where the B_{mn} are the Fourier amplitudes of magnetic field strength B in Boozer flux coordinates, and $\langle \dots \rangle_s$ denotes a weighted average (these weights distinct from the w_i above) over flux surface label s . Kink stability (χ_K^2) is computed using the Terpsichore code,⁸ and ballooning stability (χ_B^2) is computed with the COBRA code.⁹ Other contributions to χ^2 include constraints on the desired plasma aspect ratio A , volume-averaged β , rotational transform $\iota(s)$, maximum curvature of the plasma surface, and minimum plasma major radius R . Such constraints are imposed through contributions to χ^2 in the same manner as the transport and stability constraints. Any of these contributions may be switched on or off, or their relative contribution changed by changing the corresponding w_i .

III. THE DE ALGORITHM

As noted above, the DE algorithm is similar to genetic algorithms (GAs), but suited to exploration of a continuous space. Here we briefly specify the method and touch on some of its properties. The reader is referred to the fuller discussion in the original work.⁵

As indicated in Sec. I, the optimizer is initialized in generation $g=0$ with an ensemble of N_p system points $\{\mathbf{z}^i\}$ ($g=0$) distributed in some way over the D -dimensional $\mathbf{z}$ space. For generation $g+1$, a new point $\mathbf{z}^i(g+1)$ is generated from those in generation g by first using the rule

$$\mathbf{v}^i = \mathbf{z}_b(g) + (\mathbf{z}_{r_2}(g) - \mathbf{z}_{r_3}(g))F, \quad (2)$$

to produce trial vector $\mathbf{v}^i$. Here F is an input parameter, typically in the interval (0.5,1), multiplying the difference between two randomly selected system points of generation g , and setting the spread of trial points for generation $g+1$. $r_{2,3}$ are random numbers $\in [0, N_p - 1]$, and $\mathbf{z}_b(g)$ is a third element of generation g , whose choice varies with the particular DE “strategy” chosen. Other strategies arise from adding two difference vectors instead of the single one in (2). For the studies in this paper, strategy “DE/best/1” is employed,⁵ meaning that $\mathbf{z}_b(g)$ is the $\mathbf{z}$ having the lowest χ in generation g , and that only a single difference vector is employed.

Next, L of the components v_j^i of $\mathbf{v}^i$ are “mutated” to produce trial vector $\mathbf{u}^i$, similar to crossover in the GA, according to

$$u_j^i = \begin{cases} v_j^i & \text{for } j = n, n+1, \dots, (n+L-1)_D, \\ (\mathbf{z}^i)_j(g) & \text{otherwise.} \end{cases}$$

Here, n and L are integers $\in [0, D-1]$, with n the first component in the mutated sequence, randomly chosen, L the length of the mutated sequence, chosen with probability $P(L=\nu) = P_c^\nu$, with P_c the “crossover probability,” and $\langle m \rangle = m \text{ modulo } D$.

Finally, if the cost function value $\chi(\mathbf{u}^i)$ is smaller than that for $\mathbf{z}^i(g)$, $\mathbf{u}^i$ replaces that member: $\mathbf{z}^i(g+1) = \mathbf{u}^i$. In the studies discussed here, the choices $F=0.9, P_c=1$ have been made, the latter implying that no mutation is used. This has been found to produce adequate diversity in the population.

From (2), with $F \sim 1$, one sees that the trial points for generation $g+1$ will have a spread comparable to that of the current ensemble about the current-best point $\mathbf{z}_b$. If the points lie across only a single basin, the ensemble center will remain in that basin, and the search will converge rapidly to the bottom of that well. If on the other hand the $(g+1)$ th best is toward the edge of the g th ensemble, as would occur if the present ensemble lies generally on a slope, or if points at the edge have encountered a new and deeper well, the focus of the search can shift, and permit hunting in new wells. Thus, it is important that the initial ensemble be spread widely enough, in order that the evolution not be simply a convergence to the initial well bottom.

With this scheme, consistent with the guidance given⁵ for good convergence performance, we will take $N_p \approx 7D$ in the runs discussed here. $g_{\max}$ for convergence also scales with D . As a rule of thumb, we find $g_{\max} \sim 10D$. Thus, the total time for a run using N_{CPU} processors is $T = N_p \times g_{\max} \times T_1 / N_{\text{CPU}} \approx 70D^2 T_1 / N_{\text{CPU}}$, with T_1 the CPU time to evaluate a single $\chi(\mathbf{z})$. On the R10000 processors in the Origin 2000 Nirvana machines at Los Alamos National Laboratory (LANL), including a VMEC equilibrium calculation, and a Terpsichore evaluation of kink stability, we find $T_1 \approx 15$




FIG. 1. The “M1017” set of modular coils for the NCSX LI383 configuration, showing the four different types of modular coils.

minutes. Taking $D=30$, about $945 \times 10^3 / N_{\text{CPU}}$ min $= 15\,750 / N_{\text{CPU}}$ hr are thus needed to run to an optimum. Thus, having $N_{\text{CPU}} \sim 100$ (as is the case on an Origin 2000) makes such calculations accessible.

In the process of convergence in a space with $D > 1$, one finds a “dimensional contraction” in the DE ensemble (see Sec. IV), in which the ensemble converges most rapidly along the directions in $\mathbf{z}$ in which the largest variations in χ are found, followed by slower convergence along directions in which successively smaller variation occurs. The intermediate, pre-converged states of the DE ensemble thus can be of as much interest as the fully converged state, or the final state coming from a local method—instead of a single optimal $\mathbf{z}$, they typically possess a set of $\mathbf{z}^i$, with nearly optimal

χ values, but spread appreciably over $\mathbf{z}$ space. Especially since the implemented χ does not fully capture human criteria for an optimum, this property can be helpful, as will be seen.

IV. STELLOPT-DE APPLICATIONS

To illustrate some of these features of a DE search, we first show a DE study paralleling a series of LM studies done by the NCSX team to investigate the operational flexibility of the reference NCSX configuration LI383,¹⁰ using the “M1017” set of modular coils (Fig. 1). This calculation thus uses VMEC in free-boundary mode, and $\mathbf{z}$ space has dimension $D=11$. For $j=1-9$, $z_j = I_j$ are the currents in the four distinct types of modular coils, in the one set of auxiliary TF coils, and four equivalent currents for the dipole, quadrupole, hexapole, and octupole fields used to represent the PF contribution to $\mathbf{B}$. z_{10} is the central pressure p_0 and z_{11} is the enclosed toroidal flux Φ_a . These are adjusted in order to satisfy constraints for a target plasma β , and to keep the plasma limited by a specified first wall boundary. Following the heuristic rules noted in Sec. III, we take $N_p = 80 \approx 7 \times D$ for this search.

In Figs. 2(a)–2(c) are shown projections of the DE population onto two particular z_j , $j=6$ and 8, (a) for generation $g=0$, (b) for the superposition of generations $g=0-49$, and (c) for $g=49$. The choice of z_j used for plotting is arbitrary, and the appearance is similar for other pairs. The initialization randomly selects z_j values between 0.8 and 1.2 times those for LI383, producing the rectangular filled region in Fig. 2(a). Between $g=0$ and $g=49$, the ensemble spreads as the algorithm searches $\mathbf{z}$ space, as seen in Fig. 2(b). This is followed by a dimensional contraction to the nearly 1D form




FIG. 2. (a)–(c) Projections of the DE population onto z_6 and z_8 for generations (a) $g=0$, (b) $g=0-49$, and (c) $g=49$ for DE flexibility study of the LI383 configuration. (d) χ versus z_6 for the same study, for $g=49$.




FIG. 3. Scatter plot of χ_{Bmn}^2 versus χ_K^2 for an ensemble of configurations from a free-boundary run all having comparable values of total χ .

shown in Fig. 2(c). Finally, in Fig. 2(d) is shown the cost function value $\chi(\mathbf{z})$ versus z_6 for $g=49$, with the boundary line showing the locus of the cloud of values which would appear for the same superposition of generations as in Fig. 2(b). One sees that the DE evolution thus far has produced a *range* of configurations, all having cost χ comparable to the single value found by the Stellopt-LM runs, and with best value $\chi_b \approx 6.41$, slightly better than the best $\chi_{LM} \approx 6.43$ produced by the LM searches.

However, in practical terms, a configuration with lower χ^2 is not necessarily actually a better configuration, because the cost function does not perfectly express what a good configuration really is. First, in addition to physics figures of merit χ_i^2 one does regard as a measure of goodness, such as χ_{Bmn}^2 or χ_K^2 , the cost function also contains target constraints, such as the desired beta, or minimum major radius, whose relation to physics goodness is unclear, but whose contributions to χ^2 can nevertheless dominate the physics contributions. Additionally, the physics figures of merit do not always perfectly reflect the characteristic they are intended to. For example, χ_{Bmn}^2 , intended to measure confinement quality, is only a rough measure of ripple transport, and does not include axisymmetric or turbulent transport at all. And finally, in a given run, some χ_i^2 are sometimes switched off, for various reasons. These deficiencies are common to both the present LM and DE versions of Stellopt.

TABLE I. Seed configurations “a.”

Config.	L1383	PG2	C82	II75	A4k2
R [m]	1.73	1.70	1.46	1.82	1.64
I_p [kA]	150.0	150.0	200.0	120.0	230.0
RB_t [m-T]	2.05	2.02	1.66	2.16	1.85
A	4.4	4.3	3.4	4.8	4.0
$\beta_{\max}(\%)$	4.2	3.5	3.9	4.0	4.8
ι_0	0.39	0.15	0.26	0.46	0.30
ι_a	0.65	0.46	0.47	0.75	0.48
$\hat{\chi}_{Bmn}^2$	0.015	0.095	0.027	0.012	0.016




FIG. 4. Diagram showing the distances (in cm) in z space between the 10 seed configurations 1a–5a and 1b–5b used in the fixed-boundary $D=5$ restricted space.

One way the physics—versus target—constraint problem can be dealt with is illustrated in Fig. 3, which shows a scatter plot of χ_{Bmn}^2 versus χ_K^2 for an intermediate generation for a free-boundary run similar to that in Fig. 2. As in Fig. 2(d), this produces an ensemble of configurations with similar values of χ^2 , but with differing breakdowns into χ_i^2 . While all are comparably good as measured by the full χ^2 , only a few, falling within the boxed region, are regarded as satisfactory, since we are interested in configurations which are *both* quasisymmetric and kink stable. For comparison, we include here configuration “A” in Fig. 3, the optimum arrived at by the LM optimizer.

From the numerous searches with Stellopt-LM carried out by the NCSX team, a number of different fixed-boundary QA configurations with promising characteristics have been identified. One would like to know how these (and other) configurations are related to each other, how different they are from each other, and whether there might be other QA configurations in z space, as yet undiscovered, which are comparably or perhaps more attractive than those already found. In addition to the current NCSX reference configuration LI383,¹⁰ here we shall also make use of four additional QAS configurations, all having $N_f=3$ periods and fairly low aspect ratio (Table I): PG2, a QAS developed by Garabedian¹¹ to achieve good stability through a deep magnetic well, C82, an earlier reference configuration¹² constrained to fit inside the PBX vacuum vessel, II75_286b (ab-




FIG. 5. Poloidal cross sections of seed configurations 1a–5a at $\zeta=0$ and $\zeta=\pi$.




FIG. 6. Poloidal cross sections of seed configurations 1b–5b at $\zeta=0$ and $\zeta=\pi$.

breviated II75), a configuration obtained starting from LI383 with a greater target value of edge transform [$\iota_a(\text{II75}) \approx 0.75$ at $\beta=4\%$, versus $\iota_a(\text{LI383}) \approx 0.65$], and A4k2.45b4.75_b (abbr. A4k2), a configuration obtained starting from C82 but with enhanced elongation, which improves kink stability. For brevity, we refer to the resultant stellarators as configurations 1–5, in the order just given. The configurations are specified by the surface shape, hence the fixed-boundary specialization $\mathbf{z} \rightarrow \mathbf{X}$ is being used. The dimensionality needed to describe all five configurations within the $\mathbf{X}$ description is $D_0=60$.

To compare these on an equal footing, we rescale each of them to have $R_0=1.68$ m, and $\Phi_a=0.54$ Wb (the LI383 reference values). We refer to the resultant configurations as 1a–5a. These configurations were developed targeting different objectives, and thus are local optima for differing sets of weights $\{w_i\}$, and so differing χ . To make comparison among them, one must choose a single set $\{w_i\}$, for which stellarators 1a–5a will no longer be optima. For this, we choose the w_i from the prescription that each of the $w_i^2 \hat{\chi}_i^2$ contributes equally to the total χ^2 when its $\hat{\chi}_i^2$ equals an acceptably small target value. For example, the QA weight w_{Bmn} is calculated from the condition $w_{Bmn}^2 \hat{\chi}_{Bmn}^2 = 10$ when $\hat{\chi}_{Bmn}^2 = 0.015$, and the kink weight w_K is calculated from $w_K^2 \hat{\chi}_K^2 = 10$ when $\hat{\chi}_K^2 \equiv \lambda_K^2 = (-10^{-5})^2$, where λ_K is the kink eigenvalue from Terpsichore, negative for unstable modes.¹³ Accordingly, with the prescribed set $\{w_i\}$, we run Stellopt–LM on each of configurations 1a–5a, to obtain new local optima, called 1b–5b, respectively. We find that these lie fairly near the original configurations 1a–5a (Fig. 4). Here, we have introduced the simple norm $|\mathbf{z}| \equiv (\sum_j' z_j^2)^{1/2}$, where the prime on $\sum_j$ excludes from the sum the extra,




FIG. 7. Poloidal cross sections of the lowest- χ configuration at generation $g=24$ in fixed-boundary application 2.

TABLE II. Some important physics parameters for configurations “c.”

Config.	1c	2c	3c	4c	5c
A	4.4	4.4	4.0	4.4	4.2
$\beta(\%)$	4.2	4.1	4.2	4.2	4.2
ι_0	0.41	0.14	0.24	0.39	0.20
ι_a	0.65	0.67	0.65	0.65	0.65
χ^2	13.4	112	52.1	14.0	34.1
χ_K^2	0.123	42.4	2.68	0.135	0.159
χ_{Bmn}^2	8.96	38.5	14.2	8.58	13.8
χ_κ^2	2.33	2.46	11.1	2.76	8.01

nongeometric quantities such as p_0 and Φ_a introduced in Sec. II. Configurations 1a–5a and 1b–5b are shown in poloidal cross section at $\zeta=0$ and π in Figs. 5, and 6, respectively. This convergence of Stellopt–LM to five distinct minima when started from five different locations in $\mathbf{z}$ illustrates the tendency of local methods mentioned in Sec. I to become trapped in local, suboptimal minima.

As a second DE application, we now take the $D=5$ subspace spanned by these five “seed” configurations $\mathbf{z}^{1b-5b}$, and allow Stellopt–DE to search this space. The selection of the space makes use of knowledge found using the LM optimizer, but the search within this space is unbiased, i.e., the DE optimizer has no information on where these optima lie. From a best cost value $\chi_b(g=0)=31.6$ at generation $g=0$, χ_b falls to 3.16 at $g=24$ and 2.65 at $g=54$, about 87% and 73%, respectively, of $\chi(\text{LI383}_b)=3.62$ and $\chi(\text{A4k2}_b)=3.63$, the lowest two of the five seed values for the $\{w_i\}$ used. Thus, without information on the locus of the LM optima, the DE method appreciably improves on the five local minima found previously by LM plus human interaction.

As with application 1, however, a configuration with lower χ is not necessarily actually a better configuration. In this case, χ does not include a penalty for sharp corners. The result is illustrated in Fig. 7, which shows cross sections of the lowest- χ configuration at generation $g=24$ noted above. It lies quite close to A4k2_b in $\mathbf{z}$ space, only 2.9 cm away,




FIG. 8. Poloidal cross sections of seed configurations 1c–5c at (a) $\zeta=0$, (b) $\zeta=\pi$, and (c) $\zeta=\pi/2$. (d) Rotational transform ι of these configurations versus $s \equiv \Phi/\Phi_a$.




FIG. 9. (Color) Contour plot of $\chi(\mathbf{z})$ over the (z_{10}, z_{20}) plane for fixed-boundary application 3, using the set $(1c, 2c, 3c)$ of seed basis vectors.

compared with 19.2 cm from LI383_b. One notes that the tips at $\zeta=0$ have become too sharp to be realizable with practical coils (i.e., coils which are not too near the plasma).

As a third application, therefore, we investigate a similar subspace spanned by five seed configurations 1c–5c, with each obtained beginning as before with configurations 1a–5a, but now running Stellopt–LM with a penalty χ_κ^2 included for high curvature of the outer flux surface. Important elements of the χ^2 decomposition of these configurations are given in Table II, along with some other important physical parameters. The resultant LM-optimal configurations 1c–5c (Fig. 8) again lie fairly near 1a–5a, and retain their qualitative character in poloidal cross section, as one sees comparing Figs. 5 with Figs. 8(a) and 8(b). Configurations 1c and 4c lie quite near each other (unsurprising because configuration 4a derives from 1a), and these and 2c have a bullet-like shape with positive triangularity at $\zeta=\pi$, while the remaining 2 seeds, 3c and 5c, have negative triangularity and are indented around $\theta=0$. The resemblance of these two is also unsurprising, since, as noted above, 5a derives from 3a. Thus, the $\zeta=\pi$ cross section shapes suggest a division into 2 classes of configurations. A refinement of this simple picture is suggested by the ι profiles for these five stellarators shown in Fig. 8(d): Instead of 2c (=PG2_c) resembling 1c and 4c, it now resembles 3c and 5c. An overall taxonomy suggested by these two characteristics is thus 1c and 4c in class A, 2c in hybrid class B, and 3c and 5c in class C. From Table II, one observes that class B has much less kink stability than classes A and C, and from Fig. 8(b), the $\zeta=\pi$ cross section of class B also has less squareness than those of classes A and C. Thus, kink stability may be improved by increased squareness at $\zeta=\pi$, or by some more complicated combination of shaping harmonics not apparent from a simple visual appraisal of these cross sections.

Doing an unbiased DE run in the $D=5$ space spanned by configurations $\mathbf{z}^{1c-5c}$, the results are similar to those in application 2: From a best cost value $\chi_b(g=0)=42.1$ at generation $g=0$, χ_b falls to 3.53 by $g=16$, and thereafter remains quite flat, reaching 3.47 by $g=52$, about 94% and 59%, respectively, of $\chi(\text{LI383}_c)=3.66$ and $\chi(\text{A4k2}_c)$





FIG. 10. (Color) Contour plots of component parts (a) $\chi_\kappa(\mathbf{z})$ and (b) $\chi_{Bmn}(\mathbf{z})$ over the (z_{10}, z_{20}) plane for the calculation of Fig. 9.

$=5.84$. The configuration with the lowest χ for $g=52$ just mentioned lies only 1.0 cm from LI383_c and 1.9 cm from the related II75_c (the two lowest- χ c-configurations), versus 17.1 cm from A4k2_c. The χ^2 values for this run are somewhat higher than those of application 2, both because χ^2 has an additional contribution from χ_κ^2 , and because the pres-




FIG. 11. (Color) As Fig. 9, but now using the set $(1c, 2c, 5c)$ of seed basis vectors.

ence of χ_K^2 in computing the c-configurations which define the subspace of application 3 makes them more constrained than the b-configurations of application 2.

While the five seed configurations span a five-dimensional subspace of $\mathbf{z}$, as noted there are two pairs of closely related configurations among these, viz., (1c,4c) and (3c,5c). For the first pair, 1c=LI383_c is closest to the reference configuration LI383, and has lower χ for the sets $\{w_i\}$ used here. Both members of the second pair are of interest: 3c=C82_c is close to the earlier reference configuration C82, while 5c=A4k2_c has the lower χ . Thus, we now further restrict our space of study to that containing 1c,2c, and either of 3c or 5c. Choosing 3c first, the three seed vectors (1c,2c,3c) are all contained in a 2D space, spanned by

$$\mathbf{z} = \mathbf{z}^{1c} + a_1(\mathbf{z}^{2c} - \mathbf{z}^{1c}) + a_2(\mathbf{z}^{3c} - \mathbf{z}^{1c}), \quad (3)$$

for all values of coefficients $a_{1,2}$. Thus, from a 5D subspace of $\mathbf{z}$ chosen to contain much of the interesting physics uncovered by numerous studies with Stellopt-LM, we arrive at a 2D space which contains much of the physics interest of that 5D subspace, but over which visualization and physical interpretation of the variation of important quantities are far easier, our main purpose for this portion of the study.

Accordingly, in Fig. 9 is shown a contour plot of χ over the (z_{10}, z_{20}) plane, computed from a regular ensemble of $\mathbf{z}$ in the space. The increment between successive contours is $\Delta\chi = 50$. Again, the choice of z_{10} and z_{20} is somewhat arbitrary, and does not affect our conclusions. One observes a central “ridge” running roughly vertically, with peak values around $\chi \approx 300$. Configurations 1c=LI383_c and 2c=PG2_c lie in a triangle-shaped valley to the left of this ridge (χ in the range 0–50), and 3c=C82_c lies in a second valley to its right. No values are given outside the curve bounding all the contours shown, because VMEC or sometimes TERPSICHORE fails to converge there. The configurations in this region are too exotic (and thus probably less realizable) for these codes to operate successfully. Thus, while Stellopt-LM has been searching an in-principle unbounded $\mathbf{z}$ space, that space is bounded, and far smaller than one might have supposed, if one regards failure of VMEC or TERPSICHORE as a good indicator that the configuration is impractical. Moreover, across the limited range of practical configurations, one notes that χ manifests only a few maxima and minima. This 2D picture is consistent with results along 1D curves in $\mathbf{z}$ between interesting configurations reported earlier,¹³ and provides an explanation for the small number of attractive QAS classes which previous search has uncovered, included in the taxonomy just discussed.

Some insight into the physical origin of the topography of χ may be seen from the breakdown into its component parts. In Fig. 10 are shown contour plots of (a) χ_K and (b) χ_{Bmn} over the same region as Fig. 9. One notes that χ_K is by far the dominant factor [the contour spacing in Fig. 10(b) is only $\Delta\chi = 2$], with χ_{Bmn} having a significant role only in the flat triangular valley of near kink stability ($\chi_K \approx 0$), where the minimum in χ_{Bmn} near the triangle’s lower tip very nearly coincides with the position of the overall optimum LI383_c. One sees that the central ridge is due to increased

kink instability as configurations pass from the positive to negative triangularity form.

In Fig. 11 we show the projection over the (z_{10}, z_{20}) plane of a similar 2D space, but now using the set (1c,2c,5c) of seed vectors. This subspace is thus rotated slightly about the $\mathbf{z}^{2c} - \mathbf{z}^{1c}$ axis from the (1c,2c,3c) space of Figs. 9 and 10, and one again sees points 1c and 2c in a different cut of the elongated valley of near kink stability, next to the ridge of relative kink instability, with the third seed, now 5c, again appearing on the other side of this ridge right near the code convergence boundary which limits the region. The convergence region is somewhat larger than that for the (1c,2c,3c) space, and has some extra features toward the large z_{10} side. However, the generic features discussed for the (1c,2c,3c) space are true for the (1c,2c,5c) space as well, including the dominance of χ_K^2 in determining the overall topography of χ^2 .

V. DISCUSSION

We have implemented and exercised the DE algorithm in the Stellopt optimizer in both fixed- and free-boundary searches of stellarator configuration space, and found that it is, as expected, much less inclined to become trapped in local suboptimal minima than the local LM algorithm, which has been used by Stellopt until now. Stellopt-DE improved very slightly on the best value found using Stellopt-LM for free-boundary studies of operational flexibility. It made somewhat greater improvement ($\sim 30\%$) in a fixed-boundary run in a 5D restricted space containing five previously discovered interesting QAS configurations. In both cases, the intermediate generations of the DE search have produced a set of configurations which differ appreciably, but have almost the same total value of χ . As expected, a single run with the algorithm takes appreciably longer (typically a factor of 7–10) than a single LM run. However, the LM optimizer typically must be run several times, with human readjustment between runs, in order to arrive at a good optimum. For the applications studied thus far, Stellopt-DE has found configurations improving somewhat, but not dramatically, on those developed by the Stellopt-LM searches. We regard this as providing additional confidence that the LM+human approach the NCSX team has used to date to find good QAS configurations is working.

As noted, a configuration with optimal χ is not necessarily an optimal configuration in practical terms, because of a gap between the figures of merit $\hat{\chi}_i^2$ used in χ^2 and the criterion a human designer applies in judging goodness of the stellarator, and also because of the way both physics figures of merit and other target constraints are combined in χ . These are deficiencies of both the current LM and DE versions of Stellopt.

A further weakness of simply using either optimizer is that they do not provide, by themselves, much insight into why they arrive at the configurations they do, how many good configurations there may be, and what their distinguishing features are. The more global search pattern of the DE method has led us to address this, attempting to visualize the topography of χ over $\mathbf{z}$ -space. By examining a subspace

which contains a range of quite different candidate configurations found earlier, we found a taxonomy of a small number of QAS classes into which the optimization runs fall, and a map on which this taxonomy and all the configurations examined can be placed, and the relationship among them viewed. We find three principal QAS classes, designated A, B, and C, lying on the two sides of a ridge in χ , which is produced by enhanced kink instability as the configurations deform from the positive triangularity at $\zeta = \pi$ of classes A and B to the outward-indented, negative triangularity of class C. Class A is typified by LI383, the current NCSX reference design, class B by PG2, a design developed elsewhere,¹¹ and class C by A4k2 or C82, an earlier reference design.

Moreover, this map indicates that the extent Δz_j in components z_j over which one finds realizable configurations is not very large, when measured by the typical scale length L_j for χ to go from a maximum to a minimum. For the maps shown here, one has $\Delta z_j/L_j \sim 2-4$. This implies that the number of different QAS classes, as defined by the number of principal basins of χ over the realizable region of $\mathbf{z}$ space, is also not large, consistent with the small number (3) of main classes found in the present study. This picture is presently conjectural. It assumes that failure of VMEC or TERP-SICHORE is a good indicator that the configuration is too exotic to be practical. While the collection of interesting seed configurations used to form the reduced space for the mapping here represent a great deal of earlier searching, it is not clear that there are not other, possibly superior basins in some direction not accessed by optimizer runs. And the size of $\Delta z_j/L_j$ in some of those directions may not be as small as the directions examined thus far. Moreover, one would expect the precise topography of χ to change as important parameters here held constant (such as target β , or number N_f of field periods) are varied, and with these, the number and specific location of the basins may change. However, the picture provides a view which is consistent with the accumu-

lated record of searches thus far for attractive QA stellarators, and which offers the possibility of a large reduction in the uncertainty and complexity of our understanding of stellarator configuration space.

ACKNOWLEDGMENTS

The authors are grateful to S. P. Hirshman and E. Valeo for numerical advice and assistance, to A. H. Boozer and L-P. Ku for useful discussions, and to M. Zarnstorff for bringing the DE method to our attention.

This work was supported by U. S. Department of Energy Contract No. DE-AC02-76-CHO3073.

¹J. Nührenberg and R. Zille, Phys. Lett. A **114**, 129 (1986).

²J. N. Talmadge, V. Sakaguchi, F. S. B. Anderson, D. T. Anderson, and A. F. Almagri, *Proceedings of the 18th International Conference on Plasma Physics and Controlled Fusion Research*, Sorrento, Italy, October 2000 (International Atomic Energy Agency, Vienna, Austria), Paper IAEA-CN-77/EXP1/06.

³A. Reiman, G. Fu, S. Hirshman, L. Ku *et al.*, Plasma Phys. Controlled Fusion **41**, B273 (1999).

⁴W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, in *Numerical Recipes in Fortran 77* (Cambridge University Press, Cambridge, 1996), pp. 676ff.

⁵R. Storn and K. Price, Dr. Dobb's J. **22**, 18 (1997).

⁶D. E. Goldberg, *Genetic Algorithms in Search, Optimization, and Machine Learning* (Addison-Wesley, Reading, MA, 1989).

⁷S. P. Hirshman, W. I. van Rij, and P. Merkel, Comput. Phys. Commun. **43**, 143 (1986).

⁸D. V. Anderson, A. Cooper, U. Schwenn, and R. Gruber, in *Proceedings of the Joint Varenna-Lausanne International Workshop on Theory of Fusion Plasmas* (Editrice Compositori, Bologna, 1988), p. 93.

⁹R. Sanchez, S. P. Hirshman, J. C. Whitson, and A. S. Ware, J. Comput. Phys. **161**, 589 (2000).

¹⁰A. Reiman, L. Ku, D. Monticello, S. Hirshman *et al.*, Phys. Plasmas **8**, 2083 (2001).

¹¹P. R. Garabedian, Proc. Natl. Acad. Sci. U.S.A. **97**, 973 (2000).

¹²A. H. Reiman, G. Fu, S. Hirshman *et al.*, Plasma Phys. Controlled Fusion **41**, B273 (1999).

¹³H. E. Mynick and N. Pomphrey, Phys. Plasmas **7**, 4960 (2000).