SIESTA: A scalable iterative equilibrium solver for toroidal applications

S. P. Hirshman,[1] R. Sanchez,[2] and C. R. Cook[3]

[1]*Oak Ridge National Laboratory, Oak Ridge, Tennessee 37831-0117, USA*
[2]*Departamento de Fisica, Universidad Carlos III de Madrid, Leganes 28021, Spain*
[3]*University of Wisconsin at Madison, Madison, Wisconsin 53706, USA*

A new solver for rapidly obtaining magnetohydrodynamic (MHD) equilibria in toroidal systems in the presence of islands and stochastic regions is described. It is based on the Kulsrud-Kruskal MHD energy minimization principle. To carry out this minimization, small displacements are made around a convenient set of curvilinear coordinates obtained from a nearby three-dimensional equilibrium that assumes nested surfaces. Because the changes of the magnetic fields and pressure are small, corresponding to small changes in the initial magnetic and kinetic energies, solutions for the linearized perturbations can be used to rapidly and iteratively find lower energy states with magnetic islands. A physics-based preconditioner is developed to accelerate the convergence of the iterative procedure to obtain an ideal MHD equilibrium with broken magnetic surfaces (islands).
© *2011 American Institute of Physics*. [doi:10.1063/1.3597155]

I. INTRODUCTION

In the last decade, magnetohydrodynamic (MHD) equilibrium codes such as VMEC (Ref. 1) and EFIT (Ref. 2) have been the keystones of the optimization and data analysis algorithms used to design tokamaks and stellarators,[3] to help reconstruct plasma states from experimental data,[2,4] or to initialize stability (COBRA (Ref. 5), TERPSICHORE (Ref. 6), and CAS3D (Ref. 7)), neoclassical transport,[8] and extended MHD time-dependent codes (M3D (Ref. 9), NIMROD (Ref. 10), and PIXIE3D (Ref. 11)). Many of these codes rely on the underlying assumption of the existence of nested magnetic surfaces at least for the initialization phase. This assumption is also particularly crucial where fast and accurate solutions are needed in the context of design optimization or experimental reconstruction.

An efficient MHD equilibrium solver that avoids such restrictions without greatly impacting performance would open up significant new possibilities in all the aforementioned areas. For instance, it would allow targeting the control of magnetic islands and stochasticity during the configuration design phase and would extend the reconstruction capabilities to situations with islands or stochastic regions. In addition, it could also provide accurate, high-resolution magnetic fields to initialize extended MHD time-dependent codes, which could thus avoid the very slow physical island formation phase and better utilize the computational resources to simulate pending critical problems such as neoclassical tearing mode control.[12] Such a solver should also be designed to be useful for ITER-relevant scenarios.[13] This requires that such a code be highly scalable, being capable of advantageously using the thousands (or even tens of thousands) of processors needed to accommodate the very high spatial resolutions that ITER analysis would require.

In this paper, we introduce a new MHD equilibrium code called SIESTA intended to fill this need. SIESTA is a three-dimensional spectral MHD equilibrium code that does not assume the existence of nested magnetic surfaces. It allows magnetic surfaces to *tear*—changing the magnetic topology—by opening islands and/or stochastic regions to find lower energy states for the confined plasma subject to certain global constraints. Minimization of the total energy W (magnetic and internal) is used to find accessible lower energy equilibrium states for the plasma. Algorithmically, SIESTA has been designed to be highly scalable up to thousands of processors.

SIESTA is an iterative fluid equilibrium solver[14,15] which uses the linearized MHD equations to find lower energy states of the plasma starting from nearby equilibria with nested surfaces. Other iterative solvers—such as the PIES code[16]—exist which use very different solution methods than described here. SIESTA uses the nested magnetic flux surfaces computed by the VMEC (Ref. 1) code to provide both a set of quasi-polar *background* coordinates and an initial guess for the equilibrium which is used to initialize the iterative procedure. The curvilinear components of the magnetic field vector, and the scalar pressure, are all expanded in Fourier series in the background poloidal and toroidal angles. The MHD energy principle is used, together with pseudo-spectral methods to deal with the nonlinear $J \times B$ force, to find unique truncated "evolution" equations for these Fourier components. After several iterations, the evolution equations lead to a numerical MHD equilibrium state. The spectral nature of SIESTA provides a natural way to implement conservative forms for the evolution equations for the contravariant components of the magnetic field B and the plasma pressure p. It departs from the BETAS (Ref. 17) code which uses an inverse representation with an integrable Hamiltonian perturbation to describe a single magnetic island. SIESTA is not limited to such integrable (single resonant surface) situations and only uses an inverse representation for the background coordinate system.

The initial nested equilibria obtained from VMEC usually provide very close approximations to the final plasma equilibrium state. For this reason, the departures of the MHD energy W from that given by the VMEC equilibria are quite

small, even when islands (or stochastic regions) are present. Typical values of $\delta W/W \sim 10^{-4} - 10^{-6}$ depend on the rotational transform profile and the fraction of the plasma volume occupied by magnetic islands $[\delta W \sim (w/a)^4(\iota')^2]$. Here, w is the island width; a is the plasma radius; and ι' is the magnetic shear. In particular, this implies that the perturbed magnetic components and pressure will be small compared to the nested equilibrium values.

The iterative nature of SIESTA is required to solve the time-asymptotic (steady-state) nonlinear MHD force balance equation (see Eq. (2.6)). SIESTA departs from previous iterative fluid solvers by using an accurate physics-based *pre-conditioner*[18] to accelerate the convergence of the solution of the *linear* ideal MHD equations (the expansion of the nonlinear force F_{MHD} around an instantaneous plasma state *not yet* in equilibrium). The term *preconditioning* refers to transforming this linearized system into one with more favorable numerical convergence properties (smaller spectral radius or condition number) for the iterative procedure.

In SIESTA, the preconditioner is related to the *inverse* of the Hessian matrix of the full set of linearized MHD equations expressed in VMEC curvilinear coordinates. Physically, the preconditioner allows the disparity in the different MHD time scales associated with the linearized waves to be brought close to unity, thus greatly improving the convergence rate toward an equilibrium state. The efficient inversion of the Hessian is achieved by using the tri-diagonal block solver code BCYCLIC.[19] Good scalability of the calculation and application of the preconditioner have been achieved for realistic problem sizes on multiprocessor computers.

The minimum energy equilibria found by SIESTA require allowance for the departure from ideal MHD during the iterative convergence sequence. Physically, the ideal (parallel) current sheets, which form around resonant surfaces to *prevent* the formation of islands in codes such as VMEC, must be allowed to diffuse by breaking the ideal *frozen-flux* theorem for the magnetic field.[20] This is done in SIESTA by interlacing (a form of time-splitting) the preconditioned ideal iteration steps with resistive (non-ideal) steps to diffuse the magnetic field and associated resonant parallel current. In real devices, due to the large plasma temperature and consequent small resistivity, this will be a very slow physical process. SIESTA accelerates this by using an enhanced (non-physical) value for the numerical resistivity. Since the resonant parallel currents are strongly peaked around rational surfaces (at least initially), this diffusion process can be done without significantly perturbing the main (non-resonant) currents.

The paper is organized as follows. In Sec. II, the MHD energy principle is briefly reviewed. The curvilinear coordinates used in VMEC and SIESTA are described in Secs. III–V. The numerical representation for the displacement vector and MHD evolution equation in curvilinear coordinates is given in Sec. VI. Boundary conditions are described in Sec. VII. Section VIII discusses the initial surface-breaking perturbations used in SIESTA. The SIESTA iterative scheme and parallelization of the preconditioner are described in Sec. IX. Section X presents examples of SIESTA equilibria for both a tokamak and a stellarator plasma configuration. Finally, conclusions are presented in Sec. XI.

II. MHD EQUILIBRIUM ENERGY PRINCIPLE

The MHD energy of a stationary plasma (velocity v = 0) with magnetic field $\boldsymbol{B}$ and pressure p is[21]

$$W = \int \left[\frac{B^2}{2\mu_0} + \frac{p}{\gamma - 1} \right] dV. \tag{2.1}$$

Here, the integral is over the volume V of the plasma. Taking the time derivative of Eq. (2.1) and using Faraday's law for the evolution of the magnetic field

$$\frac{\partial B}{\partial t} = -\nabla \times E, \tag{2.2}$$

together with particle conservation and adiabatic law $p \sim n^\gamma$,

$$\frac{\partial p}{\partial t} = (\gamma - 1)\mathbf{v} \cdot \nabla p - \gamma \nabla \cdot (p\mathbf{v}), \tag{2.3}$$

allows us to express the change in W as follows:

$$\frac{\partial W}{\partial t} = \int [-E \cdot J + \mathbf{v} \cdot \nabla p] dV - \oint S \cdot dA,$$
$$\mu_0 J = \nabla \times B \qquad \text{(Ampere's Law)},$$
$$\mathrm{S} = \mathrm{E} \times \mathrm{B}/\mu_0 + \frac{\gamma}{\gamma - 1} p\mathbf{v} \quad \text{(Poynting+Kinetic Fluxes)}. \tag{2.4}$$

With the appropriate boundary conditions, the net flux S at the boundary can be zero and therefore ignored. Using Ohm's law $E' \equiv E + \mathbf{v} \times B = \eta J$ yields the desired MHD energy principle

$$\frac{\partial W}{\partial t} = -\int [\mathbf{v} \cdot (J \times B - \nabla p) + \eta J^2] dV. \tag{2.5}$$

If the velocity field v in Eq. (2.5) is treated as a "variational" parameter (as described below), then for sufficiently small resistivity ($\eta \to 0$), an equilibrium state is reached where W becomes quasi-stationary only when the ideal MHD force balance is satisfied

$$F \equiv J \times B - \nabla p = 0. \tag{2.6}$$

In Sec. III, these equations will be expressed in a convenient basis set of curvilinear coordinates.

III. VMEC CURVILINEAR COORDINATE SYSTEM

The set of nested magnetic surfaces from the VMEC (Ref. 1) equilibrium code is convenient to use as background coordinates for the variational calculation of the general MHD equilibrium (i.e., including islands and stochastic region) described by Eq. (2.6). (Example background coordinates are given in Sec. X). This is because the VMEC equilibrium magnetic fields will be close to the final equilibrium fields—even in the presence of islands—since only very small resonant perturbations of the magnetic field are required to "tear" the nested surfaces and change their topology. In this sense, the magnetic (and pressure) perturbations from the VMEC state should be "small." The VMEC flux coordinates $(r = \sqrt{s}, u, v)$ are defined as follows through

inverse mapping relations for the cylindrical coordinates (R, φ, Z):

$$R = \sum R_{mn}(r) \cos(mu + nv),$$
$$Z = \sum Z_{mn}(r) \sin(mu + nv),$$
$$\varphi = v. \tag{3.1}$$

Here, u is the VMEC poloidal angle and v is the VMEC (cylindrical) toroidal angle (φ). In SIESTA, the "radial" flux coordinate s in VMEC is mapped to the "polar" coordinate $r = \sqrt{s}$ to guarantee analyticity of the fields near the magnetic axis $s = 0$. For simplicity, "stellarator symmetry" has been assumed in Eq. (3.1). Retaining all trigonometric terms is routinely done in VMEC and is in principle straightforward to implement in SIESTA. The VMEC code outputs the Fourier coefficients $[R_{mn}(r), Z_{mn}(r)]$, along with those for the poloidal angle renormalization stream function $\lambda_{mn}(r)$. (Note that VMEC uses the sign convention $mu - nv$ for the arguments in Eq. (3.1).) These can be used to calculate the contravariant magnetic field components as follows:

$$\sqrt{g} \, B^u = \Phi'(\iota - \lambda_v),$$
$$\sqrt{g} \, B^v = \Phi'(1 + \lambda_u),$$
$$\sqrt{g} \, B^r = 0. \tag{3.2}$$

The Jacobian of the transformation between cylindrical and flux coordinates is denoted by $\sqrt{g} \equiv R(R'Z_u - R_u Z')$. It can be calculated from the mapping in Eq. (3.1). The quantity $2\pi\Phi(r)$ is the toroidal flux enclosed between the magnetic axis $r = 0$ and the radial coordinate surface $r = const$ and $\iota \equiv d\chi/d\Phi$ is the rotational transform, where $2\pi\chi$ is the poloidal flux. Subscripts denote angular derivatives and primes denote the radial (r) derivative.

IV. MAGNETIC ENERGY IN CURVILINEAR REPRESENTATION

In this section, we express the MHD energy principle in the *fixed* VMEC curvilinear background coordinates described in Sec. III. The MHD energy given in Eq. (2.1) may be expressed in these coordinates as

$$W = \int \left[\frac{B^i B_i}{2\mu_0} + \frac{p}{\gamma - 1} \right] |\sqrt{g}| \, dr du dv. \tag{4.1}$$

Henceforth, summation over repeated indices $i \in (r, u, v)$ is implied. There is a coordinate singularity at $r = 0$ where $\sqrt{g} = 0$. The cylindrical coordinate singularity at $R = 0$ is assumed to lie outside the plasma. Although the VMEC solution satisfies $B^r = 0$ everywhere in the plasma volume, the contravariant radial magnetic field B^r in SIESTA is in general nonzero in Eq. (4.1).

The covariant components of the magnetic field B_i in Eq. (4.1) are linearly related to the contravariant components B^i by the metric tensor elements g_{ij}

$$B_i = g_{ij} B^j,$$
$$g_{ij} \equiv R_i R_j + Z_i Z_j + \delta_{vv} R^2. \tag{4.2}$$

The subscripts on R, Z in Eq. (4.2) denote spatial partial derivatives ($R_r \equiv \partial R/\partial r$, etc.). Note that $g_{ij} = g_{ji}$ is a symmetric, positive definite tensor that may, like the Jacobian $\sqrt{g}$, be computed from the VMEC coordinate transformation in Eq. (3.1). Using Eq. (4.2), we can rewrite the variation of Eq. (4.1) as follows:

$$\delta W = \int (\mu_0^{-1} B_i \, \delta b^i + \delta P) dr du dv,$$
$$b^i \equiv \sqrt{g} B^i,$$
$$P \equiv \sqrt{g} p. \tag{4.3}$$

Note that the variation is performed keeping the background geometry fixed (that is, both g_{ij} and $\sqrt{g}$ are held constant during the variation). It is then convenient to combine the Jacobian factor to define the new set of intrinsic physical variables in Eq. (4.3). In Sec. V, the variations of the intrinsic variables will be computed explicitly.

V. VARIATIONAL PRINCIPLE IN CURVILINEAR COORDINATES

The variational equations[14] for the magnetic fields and pressure are based on the discrete time-integrated forms of Eqs. (2.2) and (2.3)

$$\delta B = -\nabla \times (E\Delta t) = \nabla \times (\xi \times B),$$
$$\delta p = (\gamma - 1)\xi \cdot \nabla p - \gamma \nabla \cdot (p\xi). \tag{5.1}$$

Here, $\xi \equiv v\Delta t$ is the perturbed displacement vector, which now will be treated as an *unconstrained* variational variable. The ideal MHD relation $E + v \times B = 0$ was used to eliminate the electric field in terms of the displacement. (Note that in order to allow breaking of magnetic surfaces, a small amount of resistivity must be applied periodically during the relaxation to an equilibrium state. This is primarily important for relaxing the large parallel current perturbations that would otherwise arise at rational surfaces where the rotational transform ι is rational.)

Equation (5.1) can be used to obtain the variations of the contravariant magnetic components required in Eq. (4.3). The following vector identity[22]:

$$\nabla \alpha \cdot \nabla \times A \equiv \nabla \cdot (A \times \nabla \alpha), \tag{5.2}$$

is useful, together with the curvilinear expression for the divergence operator

$$\nabla \cdot A = \frac{1}{\sqrt{g}} \frac{\partial a^\alpha}{\partial \alpha}, \tag{5.3}$$

where $a^\alpha \equiv \sqrt{g} A^\alpha$. Using these relations, the contravariant components of δB in Eq. (5.1) can be written as

$$\delta b^i = -\sum_{j,k} \varepsilon_{ijk} \frac{\partial(E_k \Delta t)}{\partial x_j}. \tag{5.4}$$

Here ε_{ijk} is the Levi-Civita permutation symbol. Note

$$E_j \Delta t \equiv -\sqrt{g} \, \nabla x_m \times \nabla x_n \cdot (\xi \times B)$$
$$= -\varepsilon_{jmn} \xi^m b^n. \tag{5.5}$$

Using the epsilon contraction identity $\sum_i \varepsilon_{ijk}\,\varepsilon_{imn} = \delta_{jm}\,\delta_{kn} - \delta_{jn}\,\delta_{km}$ yields

$$\delta b^i = \frac{\partial}{\partial x_j}\left(\xi^i b^j - \xi^j b^i\right). \tag{5.6}$$

Note the conservative form in Eq. (5.6)—which looks like a Cartesian expression but is valid even for curvilinear coordinates—is particularly well-suited for computing spectral components of b^i with minimal spectral pollution. It is also convenient for performing the integration by parts in the variational integrand in Eq. (4.3) required to isolate the coefficients of the independent variable ξ.

The variation of the pressure variable $P \equiv \sqrt{g}\,p$ can also be written as the sum of a conservative term and a remainder as follows:

$$\delta P = -\gamma\frac{\partial(P\xi^j)}{\partial x_j} + (\gamma - 1)\sqrt{g}\,\xi^j\frac{\partial p}{\partial x_j}. \tag{5.7}$$

Note that the first (conservative) term in Eq. (5.7) will make no contribution to the variational integral in Eq. (4.3) provided the boundary terms vanish (which they do, see the discussion in Sec. VII). However, it is necessary to retain since it *does* make a stable contribution to the Hessian (see Sec. IX A).

Inserting Eqs. (5.6) and (5.7) into the curvilinear variational principle, Eq. (4.3), and integrating the magnetic field terms by parts yields

$$\delta W = -\int \sqrt{g}\,\xi^i \cdot F_i\, dr\,du\,dv. \tag{5.8}$$

The covariant components of the ideal MHD force, Eq. (2.6), are given as follows:

$$F_i = \varepsilon_{ijk}\sqrt{g}\,J^j B^k - \frac{\partial p}{\partial x_i}. \tag{5.9}$$

Here, the contravariant components of the current J are

$$\mu_0\sqrt{g}\,J^i = \varepsilon_{ijk}\frac{\partial B_j}{\partial x_k}. \tag{5.10}$$

Once again using the epsilon contraction identity yields the following succinct expression for the force components:

$$F_i = \mu_0^{-1}B^j\left(\frac{\partial B_i}{\partial x_j} - \frac{\partial B_j}{\partial x_i}\right) - \frac{\partial p}{\partial x_i}. \tag{5.11}$$

Equation (5.11) results directly from integrating Eqs. (5.6) and (5.7) by parts in the curvilinear variational principle, Eq. (4.3).

The steepest descent method is the simplest iterative technique for minimizing the energy. It results by taking $\xi^i = P^{ij}F_j$ for some positive-definite matrix P. Inserting this into Eq. (5.8) yields

$$\delta W = -\int P^{ij}F_i F_j\, dV$$
$$\leq 0. \tag{5.12}$$

In this form, δW vanishes if, and only if, the MHD force vanishes. Equation (5.12) is an energy minimization principle in the sense that only local stable minima may be attained. The typical choice $P^{ij} = g^{ij}$ yields $||F^2||$ for the quadratic form in Eq. (5.12), but it is not the *only* possible choice. As we shall discuss, this descent method can be greatly accelerated by choosing the matrix P to be an approximate *preconditioner* that coalesces the eigenvalues of the Hessian matrix (which is the gradient of the linearized covariant forces with respect to the contravariant displacements) towards unity.

Since there may be both positive (unstable) and negative eigenvalues in the spectrum of the Hessian (except very close to the stable equilibrium state, where all eigenvalues must be negative), care must be taken when selecting the preconditioner particularly for initial states "far" from final equilibrium. To the extent that the VMEC background coordinates describe a "nearby" 3D equilibrium solution, the energy for driving these unstable modes should already be substantially reduced (compared with starting from a set of cylindrical coordinates, for instance). Nevertheless, in general, it is still necessary to shift the eigenvalue spectrum for the preconditioner in order to maintain the positive-definiteness of P required in Eq. (5.12) for convergence. As the force residual approaches zero during a descent to a stable equilibrium state, the eigenvalue shift can also be reduced towards zero. Thus, in SIESTA the preconditioner tends toward the full (stable) Hessian and a non-physical solution is avoided. This method of relating the eigenvalue shift of the preconditioner to the decreasing force residual is called pseudo-transient continuation.[23] Its implementation in SIESTA will be described in Sec IX C.

VI. NUMERICAL EXPANSION IN CURVILINEAR COORDINATES

A. Angular representation

The angular dependence of the variables appearing in Faraday's law and the pressure equation is expanded in Fourier series. For the case of stellarator symmetry, these have the following forms and parities:

$$\begin{aligned}
P &\equiv \sqrt{g}\,p = \sum p_{mn}(r)\cos(mu + nv),\\
b^r &\equiv \sqrt{g}\,B^r = \sum b^r_{mn}(r)\sin(mu + nv),\\
b^\alpha &\equiv \sqrt{g}\,B^\alpha = \sum b^\alpha_{mn}(r)\cos(mu + nv),\ \alpha \in (u, v).
\end{aligned} \tag{6.1}$$

It is straightforward to add the additional Fourier components to describe non-stellarator symmetric plasmas. The Jacobian factor $\sqrt{g}$ in Eq. (6.1) may have both angular and radial dependencies. Its appearance in the variable definitions in Eq. (6.1) minimizes the Fourier mode coupling in the evolution laws when they are written in conservative form. The sums over the Fourier modes are restricted to positive m's $(0 \leq m \leq M)$ and a symmetric range of n's $(-N \leq n \leq N)$. The set of equally-spaced collocation points used for the angular coordinates u and v are related to the M and N values by Nyquist criteria and can be extended by applying the 3/2 rule[24] to minimize aliasing effects when

evaluating nonlinear terms in the force and evolution equations.

For the stellarator-symmetric parities given in Eq. (6.1), the displacement "fluxes" have the following forms:

$$\Gamma^r \equiv \sqrt{g}\xi^r = \sum \Gamma^r_{mn}(r)\cos(mu+nv),$$

$$\Gamma^\alpha \equiv \sqrt{g}\xi^\alpha = \sum \Gamma^\alpha_{mn}(r)\sin(mu+nv), \alpha=(u,v). \quad (6.2)$$

B. Radial representation

The motivation for the radial discretization scheme is based on the observation that the total energy can be discretized to second order accuracy in the radial mesh spacing $h_r = 1/(N_r - 1)$ by using the midpoint rule to evaluate the integrand in Eq. (4.1)

$$W = h_r \sum_{j=1}^{N_r-1} \iint dudv\, w_{j+1/2} + O(h_r^2),$$

$$w_{j+1/2} = (\sqrt{g})_{j+1/2}\left[\frac{1}{2}B^2_{j+1/2} + p_{j+1/2}/(\gamma-1)\right]. \quad (6.3)$$

Here, $B^2_{j+1/2} \equiv (B^k B_k)_{j+1/2}$. Regarding the distribution of radial grid points (corresponding to indices j and $j+1/2$), $r_{j+1/2} = h_r(j-1/2)$, $j = 1, N_r - 1$ define radial "half-grid" points centered between the radial nodes ("full-grid" points) $r_j = h_r(j-1)$, $j = 1, N_r$. The notation $X_{j+\frac{1}{2}} \equiv X(r_{j+\frac{1}{2}}, u, v)$ denotes values of the quantity X on the half-radial grid.

This differencing scheme implies that the dependent variables in Eq. (6.1) appearing in the energy density w should be evaluated at half-radial points. To minimize the radial coupling of the MHD forces (described below), it is also necessary to evaluate the (independent) displacements $\Gamma^\alpha(r_j) \equiv \Gamma^\alpha_j$ in Eq. (6.2) at the *full* radial mesh points r_j. Note that $r_1 = 0$ is the coordinate axis and $r_N = 1$ is the plasma boundary.

C. Radial discretization of conservative form of evolution equations

The components of the conservative form of the ideal magnetic perturbation in Eq. (5.6) can be written as

$$\delta b^r_{j-1/2} = -\frac{\partial(e_v)_{j-1/2}}{\partial u} + \frac{\partial(e_u)_{j-1/2}}{\partial v},$$

$$\delta b^u_{j-1/2} = -\frac{\partial(e_s)_{j-1/2}}{\partial v} + (e'_v)_{j-1/2},$$

$$\delta b^v_{j-1/2} = -(e'_u)_{j-1/2} + \frac{\partial(e_s)_{j-1/2}}{\partial u}. \quad (6.4)$$

Here, $e_\alpha \equiv \Delta t E_\alpha$. Note that the u and v derivatives become algebraic in m and n when Fourier series for the electric field components are inserted in Eq. (6.4). The real-space radial operations from the full-radial grid to the half-radial grid are defined as

$$(X)_{j+1/2} \equiv \frac{1}{2}(X_{j+1} + X_j),$$

$$(X')_{j+1/2} \equiv \frac{X_{j+1} - X_j}{h_r}. \quad (6.5)$$

For the ideal magnetic perturbations, the electric field components on the full radial grid can be expressed in terms of the displacements as follows:

$$(e_r)_j = -(\Gamma^u_j[B^v]_j - \Gamma^v_j[B^u]_j),$$

$$(e_u)_j = -(\Gamma^v_j[B^s]_j - \Gamma^r_j[B^v]_j),$$

$$(e_v)_j = -(\Gamma^r_j[B^u]_j - \Gamma^u_j[B^s]_j). \quad (6.6)$$

The magnetic field components in Eq. (6.6) are computed by radially averaging the perturbed quantities in Eq. (6.4) (at interior points not including the mesh boundaries at $r=0$ and $r=1$, which must be treated as special points as described in Sec. VII) using the half-to-full mesh projection operator

$$[X]_j = \frac{1}{2}(X_{j+1/2} + X_{j-1/2}). \quad (6.7)$$

Similarly, the radial discretization of the pressure equation becomes

$$\delta P_{j+1/2} = -\gamma\left(\frac{p_{j+1}\Gamma^r_{j+1} - p_j\Gamma^r_j}{h_s} + \frac{\partial(p\Gamma^u)_{j+1/2}}{\partial u} + \frac{\partial(p\Gamma^v)_{j+1/2}}{\partial v}\right)$$

$$+ (\gamma-1)\left(p'\Gamma^r + \Gamma^u\frac{\partial p}{\partial u} + \Gamma^v\frac{\partial p}{\partial v}\right)_{j+1/2}. \quad (6.8)$$

The pressure gradient term p' on the *full* radial mesh is

$$p'_j = \frac{p_{j+1/2} - p_{j-1/2}}{h_r}. \quad (6.9)$$

This can be used to evaluate the radial pressure gradient at interior points (see Sec. VII for evaluation at $r=0$).

D. Discrete representation of the MHD forces

Inserting the discrete forms for the magnetic and pressure perturbations into Eq. (6.3) yields the following form for the full mesh covariant components of the MHD forces:

$$\delta W = -\sum_{j=1}^{N_s-1} \delta_j h_r \iint dudv(\Gamma^\alpha F_\alpha)_j,$$

$$(F_r)_j = -p'_j + (B^v K^u - B^u K^v)_j,$$

$$(F_u)_j = -\frac{\partial p_j}{\partial u} + (B^r K^v - B^v K^r)_j,$$

$$(F_v)_j = -\frac{\partial p_j}{\partial v} + (B^u K^r - B^r K^u)_j. \quad (6.10)$$

Here, $\delta_j = 1/2$, for $j=1$ and $j = N_r$ and $\delta_j = 1$, for $j = (2,\ldots,N_r-1)$. The radial contravariant components of the currents $K^\alpha \equiv \sqrt{g}J^\alpha$ are given on the full radial mesh as

$$K^r_j = \frac{\partial(B_v)_j}{\partial u} - \frac{\partial(B_u)_j}{\partial v},$$

$$K^u_j = \frac{\partial(B_r)_j}{\partial v} - \frac{(B_v)_{j+1/2} - (B_v)_{j-1/2}}{h_r},$$

$$K^v_j = \frac{(B_u)_{j+1/2} - (B_u)_{j-1/2}}{h_r} - \frac{\partial(B_r)_j}{\partial u}. \quad (6.11)$$

The angular dependence of the forces is a consequence of inserting the spectral expansions in Eqs. (6.1) and (6.2)

into δW in Eq. (6.10). A pseudo-spectral technique for the magnetic components and pressure is used to evaluate them in real space from the Fourier inverse of the evolved physics intrinsics defined in Eq. (6.1). Also, the covariant components of the current intrinsics in Eq. (6.11) are naturally filtered (as a consequence of the variational principle) over the *same* finite Fourier basis set of m and n used to compute the other intrinsic variables.

SIESTA finds the solution $F_\alpha(m, n, j) = 0$ for the three independent discrete components of the forces in Eq. (6.10), subject to suitable boundary conditions described in Sec. VII. Here, the range of Fourier indices m and n of the forces is the same used in the definition of the field and current intrinsic variables.

VII. BOUNDARY CONDITIONS

A. Edge conditions

At the plasma boundary (which is mapped to the flux surface $r = 1$), we assume a perfectly conducting wall so $B^s = \delta B^s = 0$. From Faraday's Law, Eq. (2.2), the tangential components of the electric field vanish, which in curvilinear coordinates is $E_u = E_v = 0$. Together with the ideal MHD Ohm's law, this implies $\xi^s = 0$ which makes the total energy flux S in Eq. (2.4) vanish. Equivalently, the radial integration by parts of Eq. (5.6) also vanishes in the variational principle for this choice of boundary condition. The remaining components of the displacement may be determined by evaluating the evolution equations at the radial grid point displaced by $\frac{1}{2}h_r$ from the boundary. Presently, the angular displacements are set to zero at the boundary. With this boundary condition, the radial pressure gradient is absent from Eq. (6.7) at the edge.

B. Condition at the interior coordinate singularity $r = 0$

The boundary conditions at the coordinate singularity $r = 0$ (where the Jacobian vanishes) are obtained from the physical conditions[25]

$$\frac{\partial p}{\partial u} = 0,$$
$$\frac{\partial \vec{q}}{\partial u} = 0. \tag{7.1}$$

Equation (7.1) pertains to any scalar p and any vector $\vec{q}$. The VMEC coordinates describe the coordinate axis parametrically (as a function of the toroidal angle v) as $R = R_0(v), Z = Z_0(v)$. In the neighborhood of the axis, the coordinate surfaces are asymptotically elliptical in shape, so that local displacements may be expanded as follows:

$$x \equiv R - R_0 = r \sum_n R_{1n} \cos(u + nv) + O(r^2),$$
$$y \equiv Z - Z_0 = r \sum_n Z_{1n} \sin(u + nv) + O(r^2). \tag{7.2}$$

First, consider the expansion of the scalar pressure p:

$$p(R, Z) = p(R_0, Z_0) + \left(\frac{\partial p}{\partial R}\right)_0 x + \left(\frac{\partial p}{\partial Z}\right)_0 y + O(r^2). \tag{7.3}$$

For the special case, when the magnetic axis exists and happens to coincide with the coordinate axis, then the first-order derivatives (coefficients of x and y) both vanish. However, in general, that does not occur. From Eq. (7.3), only the $m = 0$ component of p and the $m = 1$ component of dp/dr are nonzero at $r = 0$ $(j = 1)$

$$p_{j=1} = \sum_n p_{3/2}^{m=0,n} \cos(nv) + O(h_r^2),$$
$$p'_{j=1} = \sum_n \frac{2p_{3/2}^{m=1,n}}{h_r} \cos(u + nv) + O(h_r^2),$$
$$p_{j=1}^{m>1,n} = 0. \tag{7.4}$$

Now consider the asymptotic behavior for vector components. Since $\vec{q} = \sum_\alpha q^\alpha \partial \vec{x}/\partial \alpha$ and $\vec{x} = Re_R + Ze_Z$, we deduce from Eqs. (7.1) and (7.2)

$$\frac{\partial q^\varphi}{\partial u} = 0,$$
$$\frac{\partial (q^r x_r + q^\vartheta x_\vartheta)}{\partial u} = 0,$$
$$\frac{\partial (q^r y_r + q^\vartheta y_\vartheta)}{\partial u} = 0. \tag{7.5}$$

The (contravariant) toroidal component behaves like a scalar, as do the two linear combinations of the r, u components. From Eq. (7.2), note that as $r \to 0$, $x_r \approx x/r$, and $x_{\vartheta\vartheta} \approx -x$ (and similarly for y). Then, the following relations are derived (since x and x_ϑ are independent):

$$\frac{\partial q^r}{\partial u} - (rq^u) = 0,$$
$$q^r + \frac{\partial (rq^u)}{\partial u} = 0. \tag{7.6}$$

Therefore, it follows that near $r = 0$, only the $m = 1$ poloidal harmonics of the "quasi-polar" components of any vector $\vec{q}$ are non-zero and $q_{m=1}^r = -irq_{m=1}^u$ for $q_{m=1}^{\alpha=r,u} \sim \exp(iu)$.

Using these relations, it can be shown that near the coordinate axis, the $m = 1$ component of the (r, u) covariant forces are dominant and satisfy (in the limit $r \to 0$)

$$\frac{\partial F_u^{m=1}}{\partial u} = rF_r^{m=1}. \tag{7.7}$$

Note Eq. (7.7) is *similar* to Eq. (7.6), but with the component indices r, u reversed. In contrast, the $m = 0$ component of the (v) force dominates near the axis since it behaves like a scalar (similar to p).

The displacements (velocity perturbations) should be constrained during the variational descent process to produce only *physical* perturbations satisfying Eq. (7.6) near the axis. This is necessary so that the contravariant components of the magnetic field perturbations given by Faraday's law [Eq. (2.2)] will be physical and therefore satisfy Eq. (7.6) with $q \equiv B$. Since $\partial \xi^r/\partial u \to r\xi^u$ for the $m = 1$ (dominant) components of the displacement, then the ideal Ohm's law implies $\partial E_u^{m=1}/\partial u \to -rE_r^{m=1}$ and $E_v = E^0(v) + rE^1(v)\exp(iu) + O(r^2)$. These relations can be used to show that Eq. (7.6) is indeed satisfied by the contravariant components of the magnetic field perturbations when they are evolved by Eq. (6.4).

C. Natural boundary conditions at the coordinate singularity

The boundary condition $\xi^r \to -ir\xi^u$ at $r = 0$ for the $m = 1$ components of the displacement implies that the radial and poloidal variations are not independent there. For the displacement fluxes $\Gamma^\alpha = \sqrt{g}\xi^\alpha$ appearing in the evolution equations, the natural boundary conditions as $r \to 0$ are

$$\Gamma^{r,v}_{j=1} = 0,$$
$$\Gamma^u_{j=1} = \Gamma^u_{m=1,n}(r = 0). \tag{7.8}$$

Since $\Gamma^u_{m=1}$ is the only non-zero component of the flux at the origin, Eq. (7.8) is consistent with Eq. (7.6). A numerical check on these boundary conditions is provided by evaluating the ratio R for the $m = 1$ components of the magnetic field at the first half-radial grid point $r_{1/2}$

$$R = \frac{\|\partial b^r/\partial u - r_{1/2}b^u\|}{\|\partial b^r/\partial u + r_{1/2}b^u\|}. \tag{7.9}$$

According to Eq. (7.6), R should approach zero as h_r^2. During the energy minimization process, as the magnetic field perturbations are computed from Faraday's law, this scaling has indeed been verified numerically for numerous equilibria, thus confirming the validity of the natural boundary conditions given by Eq. (7.8).

VIII. INITIALIZATION: BREAKING IDEAL (NESTED) SURFACES

The initially nested surfaces in SIESTA, that arise from the VMEC background equilibrium, will be "frozen-in" by purely ideal perturbations.[20] Thus, no topological changes can occur to release magnetic energy unless "seed" island perturbations are introduced or if finite resistivity is allowed to break magnetic surfaces. Both procedures are implemented in SIESTA.

The ideal radial perturbation of the magnetic field is given by Eq. (5.6):

$$\delta b^r = \frac{\partial}{\partial u}(\xi^r b^u - \xi^u b^r) + \frac{\partial}{\partial v}(\xi^r b^v - \xi^v b^r)$$
$$\to \sqrt{g}B \cdot \nabla\xi^r \quad \text{where} \quad B^r \to 0. \tag{8.1}$$

Thus, in the initial unperturbed equilibrium state, for which $B^r = 0$, no radial perturbations will occur at rational surfaces where $B \cdot \nabla\xi^r = 0$, and, therefore, no islands can form. This is the consequence of the ideal frozen-in law. In SIESTA, small-amplitude islands are "seeded" around low order rational surfaces by adding a parallel component for the vector potential which is localized around these surfaces

$$\delta B = \nabla \times (A_{\parallel}B/|B|),$$
$$A_{\parallel} = a(r) \exp[i(mu + nv)] \tag{8.2}$$

The amplitude $a(r)$ is scanned in magnitude (for each resonance pair *m, n*) until a minimum value for the change in magnetic energy is obtained. The change in energy will be negative for small perturbations if the tearing mode is unstable.[26]

Additional tearing (topological changes) is allowed to occur in SIESTA between ideal iteration steps, by allowing the magnetic field to resistively diffuse and is described in Sec. IX D.

IX. ITERATIVE SCHEMES USED IN SIESTA

As discussed in Sec. V, for the steepest-descent iterations, the ideal MHD displacement vector is chosen to be of the form

$$\xi = PF_{MHD}. \tag{9.1}$$

(SIESTA implements a generalized form of this equation, similar to nonlinear conjugate gradients, given in Eq. (9.9)). For clarity, we use vector notation, in what follows, with the stipulation that the curvilinear components of the displacement and force vectors are actually used in SIESTA. Also, *P*—the preconditioning matrix—should be positive-definite to ensure monotonic reduction of the ideal MHD energy towards equilibrium ($F_{MHD} \to 0$) as given by Eq. (5.12).

The iteration scheme used in SIESTA is shown in the flow diagram in Fig. 1. The various parts of this diagram are described in the following subsections. In Sec. IX A, we describe the physics-based selection for the preconditioner used in SIESTA. In Sec. IX B, the efficient computation of the preconditioner is discussed. In Sec. IX C, it is shown how the preconditioner is used to iterate the ideal displacements to obtain the final equilibrium state. Application of resistivity to diffuse the resonant parallel current is described in Sec. IX D.

A. SIESTA preconditioner

The motivation for the preconditioner used in SIESTA is provided by considering Newton's method for finding the equilibrium state at iteration step $n + 1$, given a near-by state at step n

$$(F_{MHD})_{n+1} = (F_{MHD})_n + \delta J_{n+1} \times B_n + J_n \times \delta B_{n+1}$$
$$- \nabla\delta p_{n+1} = 0,$$
$$(F_{MHD})_n \equiv J_n \times B_n - \nabla p_n. \tag{9.2}$$

The perturbed field components and pressure in Eq. (9.2) are defined as *linear functionals* of the displacement vector $\Gamma = \sqrt{g}\xi$, as given by the discretized evolution equations in Sec. VI C. Thus, Eq. (9.2) can be written as

$$H_n(\Gamma_n) = -(F_{MHD})_n. \tag{9.3}$$

Here, the Hessian matrix operator H_n is related to the linearized MHD force operator[27] expressed in the VMEC curvilinear basis

$$H_n \equiv \frac{\partial F_n}{\partial \Gamma_n}. \tag{9.4}$$

Unlike the usual linearized MHD operator, however, it is *not* expanded about the equilibrium state (since $F_n \neq 0$) and, therefore, it is not self-adjoint during the convergence to equilibrium but only approaches self-adjointness as $F_{n \to N} \to 0$ for some finite value N of the iteration number.
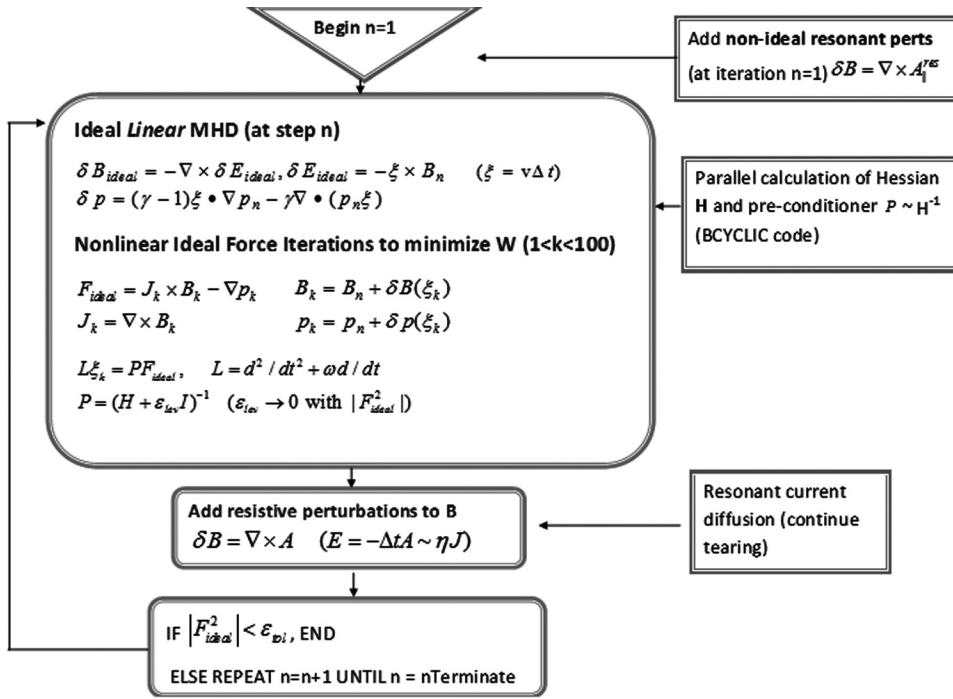
FIG. 1. Flow chart depicting the iteration scheme used in SIESTA.

However, in practice, H_n is very close to being self-adjoint $\|H - H^T\| \ll \|H + H^T\|$ as discussed below (see Eq. (9.6)).

Equation (9.3) suggests that at iteration step n, the inverse of H_n should be the preconditioner. Indeed, taking $P = H^{-1}$ (provided H is invertible) would solve the linearized ideal MHD equations exactly in one iteration step. However, since H_n may not be negative definite far from equilibrium (there may be interchange and kink instabilities leading to locally growing modes with positive eigenvalues), such a choice would have two undesirable consequences: (1) the resulting variational descent Eq. (5.12) could be violated, leading to an unphysical equilibrium state and (2) the presence of very small eigenvalues in the spectrum of H_n could lead to large displacements in the directions of the (near) null space eigenvectors. This might lead to large perturbations of the *nonlinear* ideal MHD force.

It is therefore desirable to find an approximate preconditioner based on H_n in order to assure the monotonic decrease of the magnetic energy. Referring to Eq. (5.1), note that the terms $\delta J_{n+1} \sim \mu_0^{-1} \nabla \times (\xi \times B_n)$ and the compressional part of $\delta p_{n+1} \sim -\gamma p \nabla \cdot \xi$ both make purely negative contributions to the eigenvalues of the Hessian. Furthermore, these two terms are dominant in the effective wavenumber $k \sim \nabla$, being $O(k^2)$ compared to the $O(k)$ kink and interchange terms driving (possibly) unstable modes with positive eigenvalues. A preconditioner based on this approximation will allow the rapid convergence of high spatial-frequency modes without requiring the very small time steps based on an explicit Courant-Friedricks-Lewey (CFL) numerical stability condition.

Therefore, *far from equilibrium* (when the force residual $\left| (F_{MHD})_n \right|^2$ is large), the following, negative self-adjoint approximation to the Hessian is often useful:

$$H_n^{approx}(\xi_n) \equiv \delta J_{n+1} \times B_n - \nabla \delta (p_{n+1})_{comp}$$
$$= \mu_0^{-1} \nabla \times [\nabla \times (\xi_n \times B_n)] + \gamma \nabla (p \nabla \cdot \xi_n). \quad (9.5)$$

Multiplying H_n^{approx} in Eq. (9.5) by η and integrating over the plasma volume (with the appropriate boundary conditions) shows that this approximation is indeed self-adjoint

$$\langle \eta, H_n^{approx}(\xi) \rangle = -\mu_0^{-1} \langle \delta B(\eta) \cdot \delta B(\xi) \rangle - \gamma \langle \nabla \cdot \eta \nabla \cdot \xi \rangle$$
$$= \langle \xi, H_n^{approx}(\eta) \rangle. \quad (9.6)$$

Here, brackets indicate the plasma volume integral and $\delta B(\xi) = \nabla \times (\xi \times B_n)$ is the perturbed magnetic field at iteration step n. Further, setting $\eta = \xi$ in Eq. (9.6) shows that the approximation is negative. Physically, the first term in Eq. (9.6) arises from stable Alfven waves (both compressional and field-line bending shear waves) and the second term is associated with the compressional energy of sound waves.

Note that there are zero eigenvalues for displacements proportional to the magnetic field direction (since $\nabla \cdot B_n = 0$). Other zero eigenvalues are associated with rational surfaces where $k_\parallel = 0$ (before the appearance of magnetic islands, which broaden these resonances). These must be handled before the (approximate) Hessian can be inverted. There are a number of ways to deal with this null-space of the Hessian operator. In SIESTA, we take the simple approach of adding a small diagonal element to the Hessian (either the approximate or the exact) to shift the eigenvalue spectrum towards all negative definite values (Newton's method with Hessian modification[28]). This leads to a preconditioner of the form

$$P_n = (-\tilde{H}_n + \lambda_n I)^{-1}. \quad (9.7)$$

Here, λ_n is the eigenvalue shift at iteration step n and $\tilde{H}_n = x H_n^{approx} + (1-x) H_n^{exact}$ is the effective Hessian ($0 \leq x \leq 1$). It is a blended combination of the approximate, positive-definite Hessian in Eq. (9.5) and the exact Hessian for the full linear MHD system. The blending parameter x is chosen to be 1 initially and approaches 0 as an equilibrium state is approached.

Currently, λ_0 is a user-specified parameter input to SIESTA which, due to internal scaling, is less than one (values in the range $1 \geq \lambda_0 \geq 0.001$ are typical and produce good convergence). SIESTA uses a pseudo-transient continuation method[23] to decrease this parameter as the iterations progress towards equilibrium

$$\lambda_n \sim \lambda_0 DL_n,$$
$$L_n = \sqrt{|F_{MHD}^n|^2 / |F_{MHD}^0|^2}. \quad (9.8)$$

Physically, $\lambda_n \sim \rho / \Delta t_n$ is equivalent to adding an inertial term to the force balance equation, with a large time step approaching $\Delta t_n \to \infty$ for a finite iteration count $n = N$. The reason this can occur is due to the appearance of magnetic islands that lower the magnetic energy and give rise to stable equilibrium states (broadened of the $k_{||} = 0$ resonances). Mathematically, Eqs. (9.7) and (9.8) are similar to a Levenberg-Marquardt scheme (at least when the Hessian is negative), where λ_n plays the role of the Levenberg parameter which decreases as equilibrium is approached and allows a full, quadratically convergent Newton step to occur in the final stages of the iteration.

B. Numerical computation of the SIESTA preconditioner

In SIESTA, the Hessian matrix expressed in the VMEC curvilinear basis has the form $H(\alpha, \alpha')$, where $\alpha = (j_s, m, n, t)$. Here $j_s \in (1, N_r)$ is the radial coordinate index; (m, n) represent the Fourier indices for (u, v), respectively; and $t \in (s, u, v)$ is the vector component index of the displacement. It is convenient to assemble the Fourier and type indices into a single *block* index $B = (m, n, t)$ for $m \in (0, M)$, $n \in (-N, N)$. Here, M and N are the maximum poloidal and toroidal mode numbers used in the calculation. Since the linearized MHD equations involve at most second-order spatial derivatives (see the Appendix for explicit forms for the diagonal elements of H), the SIESTA Hessian has the block-tridiagonal matrix structure shown in Fig. 2. The block index B has a linear dimension $3 \times (1 + M) \times (2N + 1)$, and there are N_r (number of radial mesh points) block rows. Based on this structure, a mesh-coloring scheme can be used to rapidly compute the elements of H (in parallel) by making three independent sweeps of the radial mesh $S_j = j + 3k, k \in 0, [N_r/3] - 1$ for $j = 1, 2, 3$. Numerically, the Hessian is computed from a finite difference version of Eq. (9.4) and can be done for all B indices at once in each sweep. In addition, each perturbation B' index can be done in parallel, so this is a very efficient numerical way to compute the Hessian.

The inversion of the augmented Hessian to compute the preconditioner defined in Eq. (9.7) has also been parallelized in the BCYCLIC (Ref. 19) code. Scalability in both the block dimension B and the block row number N_r has been achieved to give overall very good scalability.

C. Ideal iterations: Using the preconditioner

Two types of preconditioned iterations are used in SIESTA. Far from equilibrium, when the approximate Hes-
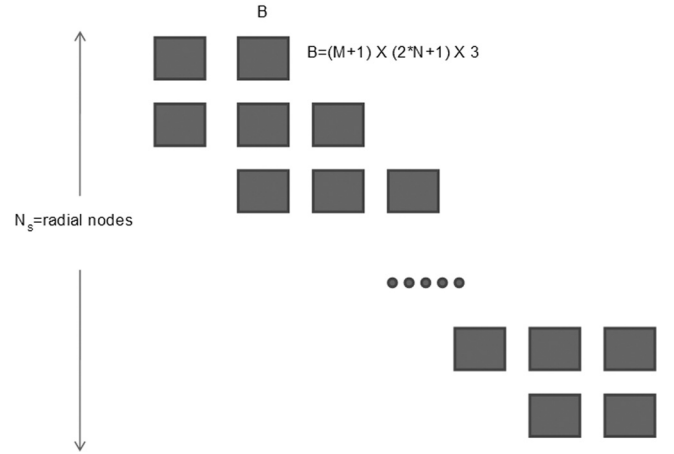
BLOCK TRIDIAGONAL STRUCTURE OF HESSIAN IN SIESTA

FIG. 2. The structure of the Hessian matrix in the curvilinear basis used in SIESTA. Each dense block consists of couplings to the 2D Fourier components of the 3 components of the displacement vector.

sian in Eq. (9.5) is used, it is convenient to use the preconditioned iteration equations

$$\frac{\partial^2 \xi}{\partial t^2} + \frac{1}{\tau} \frac{\partial \xi}{\partial t} = PF_{MHD}. \quad (9.9)$$

To the extent that the preconditioner coalesces the spectrum of the (linearized) MHD force, this iteration scheme is expected to converge very rapidly. Updates for the preconditioner are needed during the iteration sequence particularly at early times when the plasma state is far from equilibrium. Multiplying Eq. (9.9) by $\hat{H} \frac{\partial \xi}{\partial t} \equiv P^{-1} \dot{\xi}$, and integrating over the plasma volume, yields the conservation of "energy" equation

$$\frac{\partial (W_K + W)}{\partial t} = -\frac{2W_K}{\tau},$$
$$W_K \equiv \left\langle \frac{1}{2} \dot{\xi} \hat{H} \dot{\xi} \right\rangle. \quad (9.10)$$

Here, W is the plasma (potential) energy defined in Eq. (2.1) and W_K is a positive definite "kinetic" energy. This shows that the sum of the kinetic and potential energies (both positive quantities) continues to decay until the equilibrium state is achieved. Equation (9.9) represents a substantial improvement over the usual steepest descent method in Eq. (9.1) ($\tau \to 0$ limit) with an optimal nonzero choice[29] for τ.

When the force residuals are small enough (that is, when the plasma is close enough to equilibrium), the GMRES algorithm[30] can be used to solve the Newton iteration given in Eq. (9.2) with the preconditioner described in this section. SIESTA uses the CERFACS[31] "reverse communication" implementation of GMRES with the right-side preconditioner option.

D. Resistive iterations: Diffusing the resonant current

The ideal iterations described in Sec. IX C are augmented by the application of resistive diffusion to break the

ideal "frozen-in" constraint and hasten island formation and accessibility to lower energy states. The initial nested-surface VMEC state requires parallel current sheets at resonant surfaces (to *prevent* the formation of islands). Therefore, only a small resistivity is required to reduce these large resonant parallel currents. The resistivity chosen in SIESTA is a fraction of the CFL value required for numerical stability but is much larger than the physical resistivity in high-temperature plasmas. This part of the SIESTA iteration scheme is depicted in Fig. 1.

SIESTA uses a discrete form of Faraday's law, combined with Ohm's law, $E = \eta J = \eta \nabla \times B$, to diffuse the magnetic field and parallel currents

$$\delta b^\alpha = \sqrt{g}\, \nabla \cdot \eta \Delta t J \times \nabla \alpha$$
$$= \frac{\partial K_\gamma}{\partial \beta} - \frac{\partial K_\beta}{\partial \gamma},$$
$$K_\gamma \equiv (\eta \Delta t) J_\gamma. \tag{9.11}$$

Here, (α, β, γ) form a positively-ordered triplet of flux indices and the covariant components of the current $J_\beta = g_{\beta\gamma}\, J^\gamma$ are expressed in terms of b^α through Ampere's law (see Eq. (6.11)). The time step Δt is set by the CFL stability condition. Several explicit iterations of Eq. (9.11) are performed following each ideal step. The resistivity is eventually turned off when a small value of the (ideal) MHD force residual is reached to allow the ideal iteration to converge.

X. APPLICATION OF SIESTA TO THREE-DIMENSIONAL EQUILIBRIA

We will now consider two examples of 3D equilibria with islands computed using the SIESTA code. Neither of these cases is sufficiently large in terms of radial points nor Fourier modes to definitively demonstrate the parallel scaling of SIESTA (a future publication will present more detailed analysis of the scaling properties of SIESTA), but they demonstrate the new physical features and computational speed of the code. Both cases were computed on an 8-processor (2 quad-core) Linux machine and required between 3 and 5 minutes to achieve convergence for which the squared RMS normalized force residuals decreased below 10^{-20}.

A. Doublet-IIID (DIII-D) Tokamak

The equilibrium examined here is a DIII-D counter-ECCD (electron cyclotron current drive) discharge[32] (shot #133221, t = 0.3775 s). The axisymmetric (no islands) equilibrium flux surfaces equally spaced in $r \sim \sqrt{\Phi}$ (Φ is the toroidal flux) computed with the VMEC (Ref. 1) code are shown in Fig. 3. (The actual equilibrium is slightly vertically asymmetric, but it was symmetrized here so SIESTA could treat it.) The q (safety factor) and pressure profiles were computed with the EFIT (Ref. 2) equilibrium and data analysis code and are shown in Fig. 4. This equilibrium corresponds to $\langle \beta \rangle \sim 1.84\%$ and is ideally Mercier unstable in the region interior to the $q = 1$ surface.

SIESTA was run for this case with $N_s = 101$ radial surfaces, M = 12 poloidal modes, and toroidal modes $3 \leq N \leq 3$. This corresponds to $\sim 27\,500$ simultaneous nonlinear
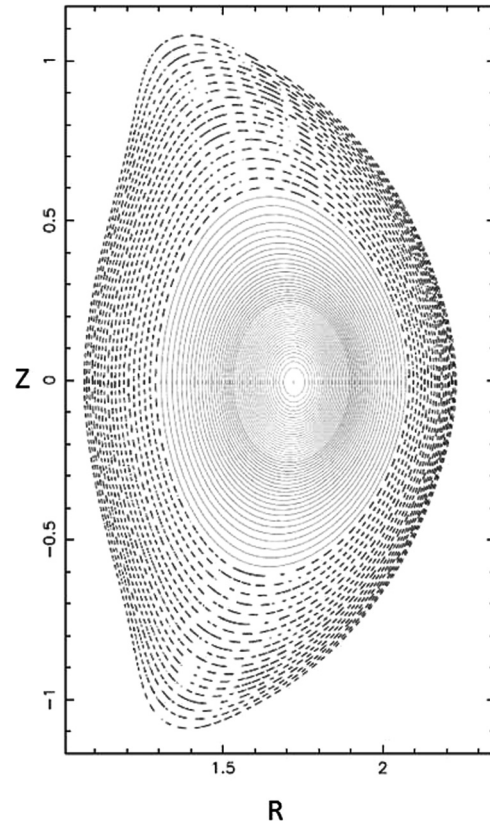

FIG. 3. DIII-D initial (symmetric) flux contours for shot #133221 described in text.

coupled differential equations. The converged SIESTA equilibrium, which allows for the formation of an island at the unstable $q = 1$ surface (thus lowering the total energy), is shown for two toroidal cross sections ($v = 0°$ and $v = 180°$) in Figs. 5 and 6. Figure 5 shows the pressure contours in real (R-Z) space, while Fig. 6 depicts the same equilibria in flux space (s = r, u). Very small initial resonant magnetic perturbations ($|B^r_{m=1-2,n=1}| < 10^{-4} B^v_{0,0}$) were applied at the $q = 1$ and $q = 2$ surfaces. The tearing unstable $q = 1$ island dominates in the saturated equilibrium state, although a small remnant of the $q = 2$ island persists (Fig. 6)
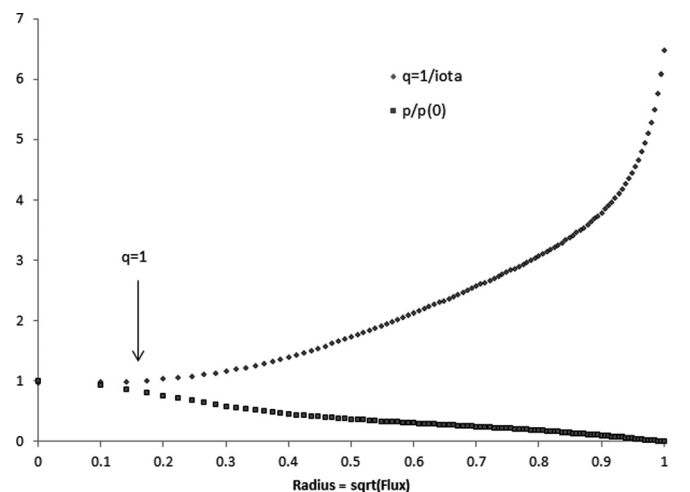

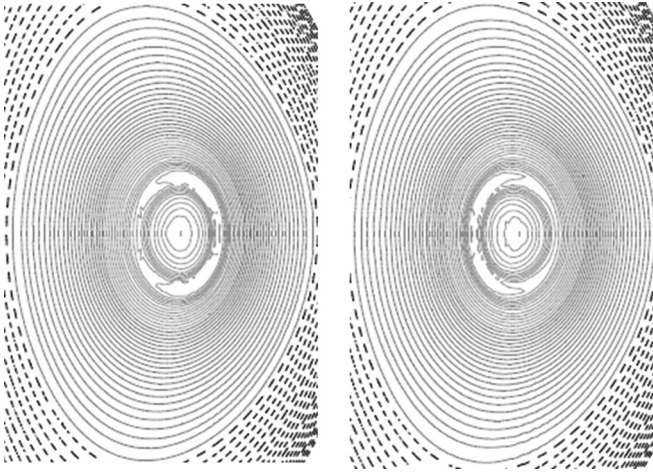FIG. 4. Normalized pressure $p/p(0)$ and safety factor (q) vs minor radius r.

FIG. 5. Pressure surfaces (excluding some unperturbed surfaces) showing the island formation at the original q = 1 surface. The symmetry planes at toroidal angles $v = 0°$ (left) and $v = 180°$ (right) are displayed. The scale is the same as in Fig. 3.

B. Classical low-beta stellarator (no net current)

The model for a (nearly) classical stellarator[33] considered here has $N_p = 3$ toroidal field periods. The fixed boundary for this stellarator is given by
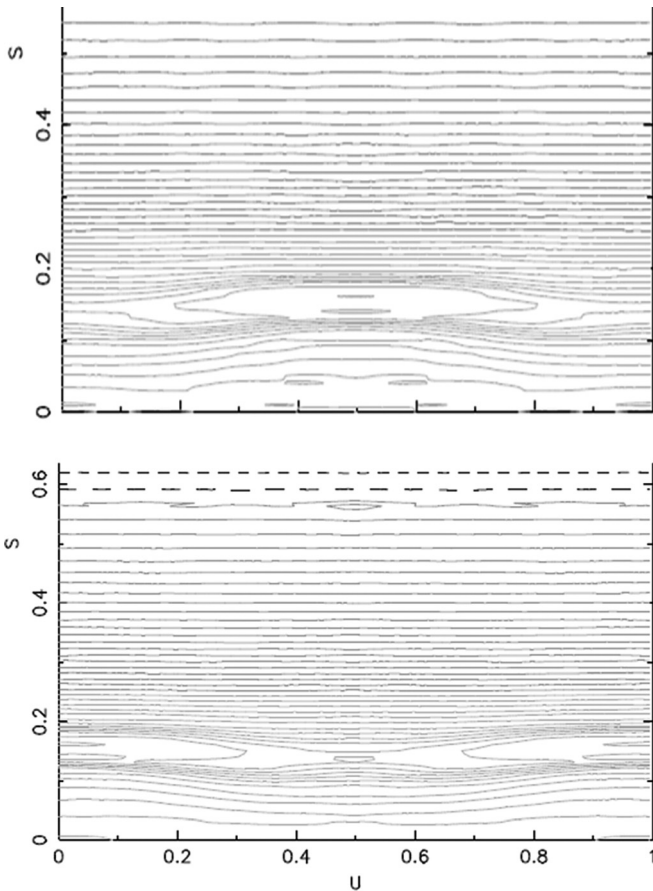

FIG. 6. Equilibrium pressure contours depicted in flux (s-u) space. Here, $s = r$ is the square root of the normalized toroidal flux and u is the poloidal angle, normed to 2π. The single period magnetic island at $s = 0.15$ (q = 1 resonance) is evident. The bulging of the contours around s = 0 indicates a small toroidally-varying shift of the original single magnetic axis. The top figure is at $v = 0$ and the bottom is at $v = 180°$. Note the phase shift of π between the two cross sections.

$$R = 2.90 + \cos u - 0.51 \cos(u + N_p v)$$
$$- 0.01[\cos(4u + N_p v) + \cos(6u + N_p v)],$$
$$Z = \sin u + 0.51 \sin(u + N_p v)$$
$$+ 0.01[\sin(4u + N_p v) - \sin(6u + N_p v)]. \quad (10.1)$$

The dominant helical terms (~ 0.51) produce a purely elliptical stellarator boundary, while the remaining terms (~ 0.01) represent small helical deformations. The unperturbed pressure surfaces from a VMEC equilibria are shown in Fig. 7 for the two symmetric cross-sections $v = 0, 2\pi/N_p$. These form the nested background coordinates for the SIESTA equilibrium calculation as previously described. The equilibrium considered is low pressure, $\langle \beta \rangle \sim 10^{-4}$ (very close to a vacuum). The helical twist of the boundary produces an iota profile of the form shown in Fig. 8. The location of the two low-order resonant surfaces, where $\iota/N = 1/m$ (for m = 6 and m = 5) and where islands are
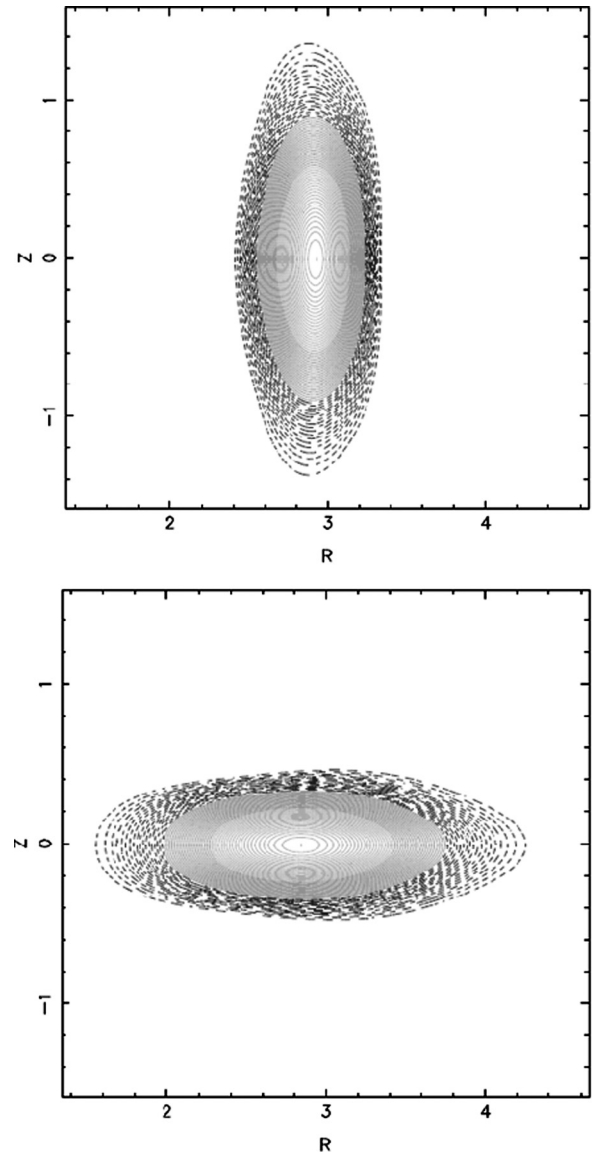

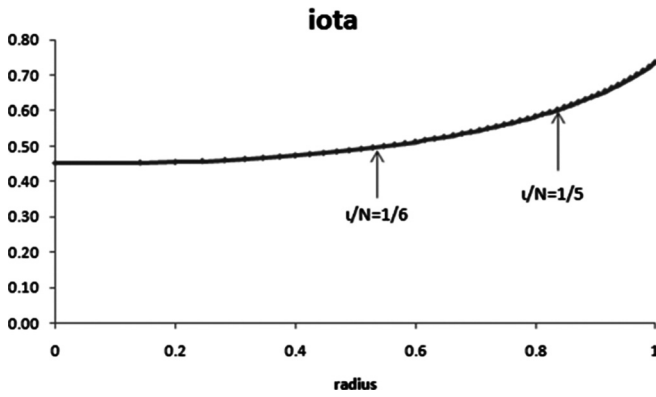FIG. 7. Initial VMEC pressure surfaces for low-beta model stellarator configuration at two toroidal symmetry planes.

FIG. 8. Iota profile for the stellarator configuration shown in Fig. 7. The two low order resonances (m = 6 and m = 5) are indicated.

expected to break the nested surface topology, are indicated by arrows.

SIESTA calculations were performed for this equilibrium using 101 radial surfaces, $M \sim 12\text{-}24$ poloidal modes and $|N|/N_p \sim 3\text{–}5$ toroidal modes. Figure 9 shows a Poincaré puncture plot where the intersection of different magnetic field lines with the toroidal plane $v = 0$ are plotted as points. The magnetic field lines were traced in flux space using the "2D orbit" equations

$$\frac{dr}{dv} = \frac{B^r}{B^v},$$
$$\frac{du}{dv} = \frac{B^u}{B^v}. \tag{10.2}$$

It was assumed that the contravariant toroidal field component B^v never reverses sign in the plasma so that the toroidal angle was used as a "time-like" coordinate. In the absence of islands (Fig. 7), the magnetic surfaces would appear as straight horizontal lines in this flux space plot. On this radial scale, the puncture plots are the same for the range of modes considered. Note the island formation at the two rational surfaces $r \sim 0.54$ and $r \sim 0.82$ with the anticipated poloidal
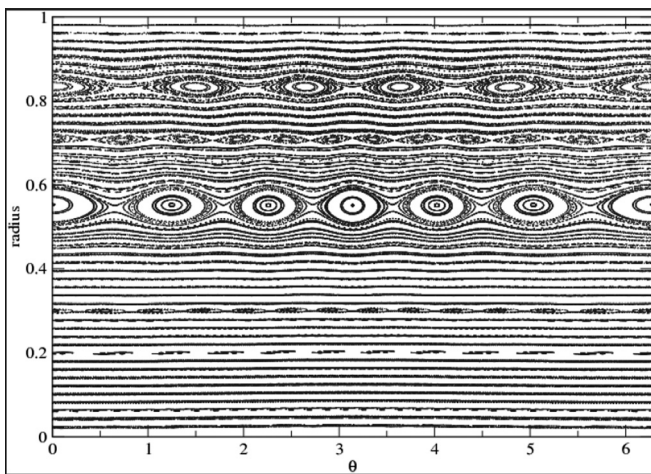
FIG. 9. Puncture (Poincare) plot in flux space, showing the intersection of magnetic field lines with the v = 0 plane for the model stellarator. The two island chains at the resonant surfaces r = 0.54 and r = 0.82 are visible.
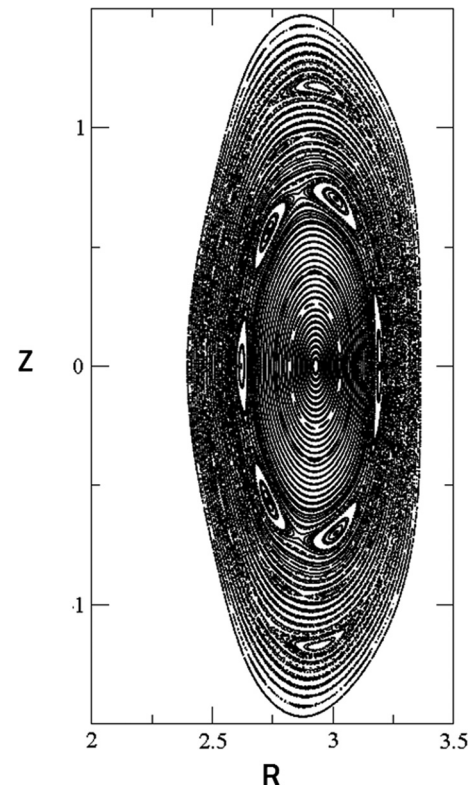

FIG. 10. The same configuration as shown in Fig. 9 but depicted in real (R, Z) coordinate space.

frequencies $m = 6$ and $m = 5$, respectively. Figure 10 depicts the surfaces in real cylindrical coordinates in the toroidal plane $v = 0$ and agrees well with results obtained from other equilibrium codes.[33]

XI. CONCLUSIONS

We have described the theory and development of a new ideal MHD 3D equilibrium solver, capable of dealing with magnetic islands and stochastic regions in a fast, accurate and scalable manner. SIESTA is particularly suited for applications that require not only accuracy but also speed of evaluation as well. These include both experimental 3D equilibrium reconstruction and stellarator design, for which SIESTA opens up new possibilities. Due to its high degree of parallelism, SIESTA could also help to calculate MHD equilibria at very high spatial resolutions, such as those that are required for the investigation of Neoclassical Tearing Mode (NTM) control at realistic ITER-relevant conditions.

APPENDIX: ANALYTIC EXPRESSIONS FOR DIAGONAL CONTRIBUTIONS TO THE LINEARIZED MHD FORCES

The leading order (in the mesh spacing and M, N) *diagonal* components of the linearized forces $F_{1\alpha} \sim L_\alpha \Gamma^\alpha$ (L_α is a second order spatial operator) are given in the limit where the radial component B^r is small as:

$$F_{1r} = B_0^u \frac{\partial}{\partial r}\left[\frac{g_{uu}}{\sqrt{g}}\frac{\partial}{\partial r}\left(B_0^u \Gamma^r\right) + \frac{g_{uv}}{\sqrt{g}}\frac{\partial}{\partial r}\left(B_0^v \Gamma^r\right)\right]$$
$$+ B_0^v \frac{\partial}{\partial r}\left[\frac{g_{vv}}{\sqrt{g}}\frac{\partial}{\partial r}\left(B_0^v \Gamma^r\right) + \frac{g_{vu}}{\sqrt{g}}\frac{\partial}{\partial r}\left(B_0^u \Gamma^r\right)\right], \quad (A1)$$

$$F_{1u} = B_0^v\left[\frac{\partial}{\partial u}\left\{\frac{g_{vv}}{\sqrt{g}}\frac{\partial}{\partial u}\left(B_0^v \Gamma^u\right) - \frac{g_{vu}}{\sqrt{g}}\frac{\partial}{\partial v}\left(B_0^v \Gamma^u\right)\right\}\right.$$
$$\left. + \frac{\partial}{\partial v}\left\{\frac{g_{uu}}{\sqrt{g}}\frac{\partial}{\partial v}\left(B_0^v \Gamma^u\right) - \frac{g_{uv}}{\sqrt{g}}\frac{\partial}{\partial u}\left(B_0^v \Gamma^u\right)\right\}\right], \quad (A2)$$

$$F_{1v} = B_0^u\left[\frac{\partial}{\partial u}\left\{\frac{g_{vv}}{\sqrt{g}}\frac{\partial}{\partial u}\left(B_0^u \Gamma^v\right) - \frac{g_{vu}}{\sqrt{g}}\frac{\partial}{\partial v}\left(B_0^u \Gamma^v\right)\right\}\right.$$
$$\left. + \frac{\partial}{\partial v}\left\{\frac{g_{uu}}{\sqrt{g}}\frac{\partial}{\partial v}\left(B_0^u \Gamma^v\right) - \frac{g_{uv}}{\sqrt{g}}\frac{\partial}{\partial u}\left(B_0^u \Gamma^v\right)\right\}\right]. \quad (A3)$$

These are useful for estimating scaling of an approximate tri-diagonal preconditioner, which can be used to effect initial reductions of the MHD forces before islands form and the full block preconditioner is applied.

[1] S. P. Hirshman and J. C. Whitson, Phys. Fluids **26**(12), 3553 (1983).
[2] L. L. Lao, H. St. John, R. D. Stambaugh, A. G. Kellman and W. Pfeiffer, Nucl. Fusion **25**, 1611 (1985).
[3] D. A. Spong, S. P. Hirshman, L. A. Berry, J. F. Lyon, R. H. Fowler, D. J. Strickler, M. J. Cole, B. N. Nelson, D. E. Williamson, A. S. Ware, D. Alban, R. Sanchez, G. Y. Fu, D. A. Monticello, W. H. Miner, and P. M. Valanju, Nucl. Fusion **41**, 711 (2001).
[4] J. D. Hanson, S. P. Hirshman. S. F. Knowlton, L. L. Lao, E. A. Lazarus, and J. M. Shields, Nucl. Fusion **49**, 075031 (2009).
[5] R. Sanchez, S. P. Hirshman, J. C. Whitson, and A. S. Ware, J. Comput. Phys. **161**, 576 (2000).
[6] D. V. Anderson, W. A. Cooper, R. Gruber, S. Merazzi, and U. Schwenn, Sci. Comput. Supercomputer **II**, 159 (1990).
[7] C. Nuhrenberg, Phys. Plasmas **6**, 137 (1999).
[8] W. I. van Rij and S. P. Hirshman, Phys. Fluids B **1**, 563 (1989).
[9] W. Park, E. V. Belova, G. Y. Fu, X. Z. Tang, H. R. Strauss, and L. E. Sugiyama, Phys. Plasmas **6**, 1796 (1999).
[10] C. R. Sovinec, A. H. Glasser, T. A. Gianakon, D. C. Barnes, R. A. Nebel, S. E. Kruger, D. D. Schnack, S. J. Plimpton, A. Tarditi, and M. S. Chi, J. Comput. Phys. **195**, 355 (2004).
[11] L. Chacón, Phys. Plasmas **15**, 056103 (2008).
[12] R. Carrera, R. D. Hazeltine, and M. Kotschenreuther, Phys. Fluids **29**, 899 (1986).
[13] "Progress in the ITER Physics Basis", Nucl. Fusion **47**, No.6, S1-S414 (2007). Also see http://www.efda.org/the_iter_project.
[14] R. Chodura and A. Schlüter, J. Comput. Phys. **41**, 68 (1981).
[15] K. Harafuji, T. Hayashi, and T. Sato, J. Comput. Phys. **81**, 169 (1989).
[16] A. H. Reiman and H. S. Greenside, Comput. Phys. Commun. **43**, 157 (1986).
[17] O. Betancourt, Commun. Pure Appl. Math. **XLI**, 551 (1988).
[18] J. Dongarra, I. S. Duff, D. C. Sorensen, and H. A. van der Vorst, *Numerical Linear Algebra on High-Performance Computers* (SIAM Publications, Philadelphia, PA, 1998).
[19] S. P. Hirshman, K. S. Perumalla, V. E. Lynch, and R. Sanchez, J. Comput. Phys. **229**, 6392 (2010).
[20] H. Alfven. J. Geophys. Res. **81**, 4019, (1976).
[21] M. D. Kruskal and R. M. Kulsrud, Phys. Fluids **1**, 265 (1958).
[22] D. L. Book, *NRL Plasma Formulary* (Naval Research Laboratory, Washington, DC, 1980).
[23] A. T. Kelley and D. E. Keyes, SIAM (Soc. Ind. Appl. Math.) J. Numer. Anal. **35**(2), 508 (1998).
[24] S. A. Orszag, J. Atmos. Sci. **28**, 1074 (1971).
[25] A. Y. Aydemir and D. C. Barnes, J. Comput. Phys. **53**, 110 (1984).
[26] A. H. Boozer, Phys. Fluids **27**, 2055 (1984)
[27] B. Bernstein, E. A. Frieman, M. D. Kruskal, and R. M. Kulsrud, Proc. Roy. Soc. London, Ser. A **244**, 17 (1958).
[28] J. Nocedal and S. J. Wright, *Numerical Optimization*, Springer, second ed. (2006) New York, NY.
[29] P. Garabedian, Math. Tables Aids Comput. **10**, 183 (1956).
[30] Y. Saad and M. H. Schultz, SIAM (Soc. Ind. Appl. Math.) J. Sci. Stat. Comput. **7**, 856 (1986).
[31] V. Frayssé, L. Giraud, S. Gratton, and J. Langou, ACM Trans. Math. Softw. **31**(2), 228 (2005).
[32] J. L. Luxon, Nucl. Fusion **42**, 614 (2002).
[33] A. Reiman, D. Monticello, and H. Strauss, private communication (2010).