Journal of Computational Physics 297 (2015) 427–441

Contents lists available at ScienceDirect

Journal of Computational Physics

www.elsevier.com/locate/jcp

Coupling extended magnetohydrodynamic fluid codes with radiofrequency ray tracing codes for fusion modeling

Thomas G. Jenkins [a,*], Eric D. Held [b]

[a] Tech-X Corporation, 5621 Arapahoe Avenue, Boulder, CO, 80303, USA
[b] Utah State University, Logan, UT, 84322, USA

ARTICLE INFO

Article history:
Received 10 December 2014
Received in revised form 11 May 2015
Accepted 19 May 2015
Available online 28 May 2015

Keywords:
Computational geometry
Ray tracing
Finite element methods
Pseudospectral methods
Magnetohydrodynamics
Radiofrequency waves

ABSTRACT

Neoclassical tearing modes are macroscopic ($L \sim 1$ m) instabilities in magnetic fusion experiments; if unchecked, these modes degrade plasma performance and may catastrophically destroy plasma confinement by inducing a disruption. Fortunately, the use of properly tuned and directed radiofrequency waves ($\lambda \sim 1$ mm) can eliminate these modes. Numerical modeling of this difficult multiscale problem requires the integration of separate mathematical models for each length and time scale (Jenkins and Kruger, 2012 [21]); the extended MHD model captures macroscopic plasma evolution while the RF model tracks the flow and deposition of injected RF power through the evolving plasma profiles. The scale separation enables use of the eikonal (ray-tracing) approximation to model the RF wave propagation. In this work we demonstrate a technique, based on methods of computational geometry, for mapping the ensuing RF data (associated with discrete ray trajectories) onto the finite-element/pseudospectral grid that is used to model the extended MHD physics. In the new representation, the RF data can then be used to construct source terms in the equations of the extended MHD model, enabling quantitative modeling of RF-induced tearing mode stabilization. Though our specific implementation uses the NIMROD extended MHD (Sovinec et al., 2004 [22]) and GENRAY RF (Smirnov et al., 1994 [23]) codes, the approach presented can be applied more generally to any code coupling requiring the mapping of ray tracing data onto Eulerian grids.

1. Introduction

Fusion plasmas in tokamaks – toroidal devices in which plasma is confined by external and self-generated magnetic fields wrapped about the major and minor axes of the torus – are susceptible to instabilities known as tearing modes (TMs) [1,2]. These modes alter the topology of the confining magnetic fields to produce helical magnetic islands that can grow to macroscopic scales (of order ~10% of the toroidal minor radius) before the TM saturates due to nonlinear effects. Tearing modes reduce density and temperature in the plasma core [3,4] as the distortion of magnetic field lines imparts an increased radial component to the rapid transport of heat and particles along these field lines. In the worst case, low-helicity TMs may give rise to rapid termination of the plasma discharge [5–7] which can damage the material walls. The development of effective strategies to suppress these modes or mitigate their effects is thus required for the next generation of fusion

* Corresponding author.
E-mail address: tgjenkins@txcorp.com (T.G. Jenkins).

http://dx.doi.org/10.1016/j.jcp.2015.05.035
0021-9991/

experiments such as the ITER tokamak [8], whose stored energy will exceed that of currently operating machines by at least an order of magnitude.

Because the initial island growth is brought about by changes to the plasma current profile, an approach to TM mitigation using electron cyclotron current drive (ECCD) has been developed. In this approach, radiofrequency (RF) waves are injected into the plasma. The wave properties are carefully chosen so as to resonate with the cyclotron motion of plasma electrons at or near the island center. Qualitatively, localized currents thus induced in the immediate neighborhood of the island counteract the TM's alteration of the plasma current profile, such that island growth can be arrested and even reversed. ECCD-based TM stabilization strategies have been highly successful in multiple experiments [9–12], and sophisticated control strategies which rely on this approach [12–16] have demonstrated excellent capability to control TMs in high-beta, fusion-relevant regimes. Nevertheless, the potential severity and expense of TM-induced disruptions in large fusion devices such as ITER suggest the value of a more quantitative understanding of the physics through predictive computational modeling.

Recent theoretical work by Hegna and Callen [17] and Ramos [18,19] (hereafter, HCR) outlines a first-principles derivation of equations that can be used to simulate both TMs and the effects of externally imposed ECCD within the formalism of extended magnetohydrodynamics (extended MHD). The HCR approach, which is valid when distribution functions of the plasma species are nearly local Maxwellians, employs a multiscale ordering wherein the spatial and temporal scales of the RF fields are assumed to be much shorter than those corresponding to fluid quantities. One may thus calculate the propagation of RF fields through effectively static fluid density and temperature profiles and feed their effects back into the profile evolution equations. More specifically, fluid equations for density, momentum, and energy are supplemented by a drift-kinetic equation ordered appropriately for the low-collisionality, high-beta regime of interest. Closures for the fluid equations are obtained from solutions to this drift-kinetic equation; they include the neoclassical bootstrap currents important for TM growth in high-beta regimes. Independently, equations for the linear propagation of RF waves through the background profiles can be solved by ray tracing methods [20]. Finally, the RF-induced perturbations can be combined to form source terms which enter quasilinearly on the fluid scales, acting as sources of momentum and energy in the fluid equations and as thermodynamic drives in the drift-kinetic equation. Analytic forms for these sources, suitable for computational implementation, have recently been derived [21].

The present work presents solutions to a number of unique challenges encountered in developing a computational implementation of the coupled RF/extended MHD equations for ECCD-based NTM stabilization. Primarily, these are issues relating to spatial resolution and convergence. They arise in considering how datasets sampled along discrete ray trajectories can be transferred to a continuous fluid representation that incorporates finite-element and/or spectral basis functions. We show that physical quantities sampled along individual ray trajectories must be combined with geometric data of the collective ray bundle to carry out this transformation. We also discuss methods by which the latter data can be calculated. While the techniques we present are generally applicable whenever continuous representations of discrete ray data are needed, this work uses the NIMROD [22] and GENRAY [23] codes. NIMROD solves the extended MHD equations using a mixed spatial representation. In toroidal geometry, 2D high-order finite elements comprise the poloidal plane, while a discrete Fourier series forms the basis in the toroidal direction. A mixed implicit/semi-implicit leapfrog algorithm is used for the time advance. NIMROD has been used to investigate many problems of relevance to magnetic fusion, including disruption dynamics [24], energetic particle stabilization effects [25], ELM onset thresholds [26], runaway electron confinement [27], tearing mode dynamics [28], and MHD closures [29]. GENRAY calculates the propagation trajectories and power dissipation associated with the passage of electromagnetic waves through a plasma. It has likewise seen considerable use in the fusion community, including the modeling of experimental lower hybrid [30,31] scenarios and the calculation of power deposition/current drive associated with injected electron Bernstein waves [32,33].

In Section 2 of this work, we present the underlying equations for the fluid and RF physics, and discuss their solutions in the NIMROD and GENRAY codes. In Section 3, we then explain the computational geometry techniques that are necessary to convert discrete RF solutions (valid only along individual rays) to the continuous quasilinear diffusion terms which drive the RF-induced fluid motion. Section 4 addresses the interpolation of the quasilinear diffusion data onto the NIMROD basis, treating both pseudospectral and finite element representations of this data. Finally, Section 5 presents some basic results of the computational simulations, summarizes the key aspects of this work, and discusses its relevance to future research.

2. Fluid and RF equations

2.1. Fluid equations and their NIMROD implementation

The extended MHD equations used in this work were derived formally in Ref. [21] for an ion–electron plasma. They are presented here in their quasineutral ($n_i \sim n_e \equiv n_s$) forms; the s subscript appearing on the density n_s and other forthcoming quantities indicates variation only on the slower, longer spatiotemporal scales of extended MHD (in contrast to the short-wavelength, rapid variation of the injected RF fields). In addition to a continuity equation in which the RF plays no role [see Eqs. (22), (73), and (74) of Ref. [21]], one obtains equations for the momentum and temperature evolution of each species

$$m_\alpha n_s \left(\frac{\partial \mathbf{V}_{\alpha s}}{\partial t} + (\mathbf{V}_{\alpha s} \cdot \nabla)\mathbf{V}_{\alpha s} \right) = -\nabla(n_s T_{\alpha s}) - \nabla \cdot \mathbf{\Pi}_\alpha + q_\alpha n_s [\mathbf{E}_{0s} + \mathbf{V}_{\alpha s} \times \mathbf{B}_{0s}] + \mathbf{R}_\alpha + \mathbf{k}_{rs} H_\alpha \delta_{\alpha,e}, \tag{1}$$

$$\frac{3}{2}n_s\left(\frac{\partial T_{\alpha s}}{\partial t}+(\mathbf{V}_{\alpha s}\cdot\nabla)T_{\alpha s}\right)+n_sT_{\alpha s}\nabla\cdot\mathbf{V}_{\alpha s}=-\nabla\cdot\mathbf{q}_\alpha-\boldsymbol{\Pi}_\alpha:\nabla\mathbf{V}_{\alpha s}+Q_\alpha+\omega H_\alpha\delta_{\alpha,e};\tag{2}$$

here q_α and m_α are the species charge and mass. The density, velocity, and temperature $(n_s,\mathbf{V}_{\alpha s},T_{\alpha s})$ of species α evolve in response to collisional momentum and energy transfer $(\mathbf{R}_\alpha,Q_\alpha)$, extended MHD-scale electric and magnetic fields $(\mathbf{E}_{0s},\mathbf{B}_{0s})$, stress tensors and heat fluxes $(\boldsymbol{\Pi}_\alpha,\mathbf{q}_\alpha)$ associated with the extended MHD closures, and quasilinear terms $(\propto H_\alpha)$ induced by the RF. Consistent with the HCR formalism [17], these quasilinear terms are valid only for electrons and are thus marked by the Kronecker delta $\delta_{\alpha,e}$. Quasineutral Maxwell equations [see Eqs. (14)–(15), and subsequent discussion, in Ref. [21]] determine the self-consistent response of the extended MHD-scale electromagnetic fields. Apart from the quasilinear terms and the details of the closures, Eqs. (1)–(2) and the other equations described (not explicitly stated as they are not of direct relevance to this work) are standard fluid equations covered by introductory texts [34]. We will not discuss the closure calculations or their associated drift-kinetic equation here (see Refs. [18,19,35] for a rigorous treatment of these issues). Rather, our focus is on computational representations of the quasilinear terms which appear as momentum and energy sources in the fluid equations (1)–(2). In these expressions, ω and $\mathbf{k}_{rs}$ are the frequency and real component of the RF wavevector, $\mathbf{k}_s\equiv\mathbf{k}_{rs}+i\mathbf{k}_{is}$, and $H_e=H_\alpha\delta_{\alpha,e}$ takes the form [21]

$$\begin{aligned}H_e=&\frac{\epsilon_0}{4}\sum_{n=-\infty}^{\infty}\frac{\omega_{pes}^2}{\omega^2}\xi_{0s}\sqrt{\pi}\exp(-\xi_{ns}^2-\lambda_{es})\\&\times\left[[I_n(\lambda_{es})-I_{n+1}(\lambda_{es})]\left(2\lambda_{es}|E_{ys}|^2+n\left|E_{xs}-iE_{ys}+\frac{\sqrt{\lambda_{es}}(\omega-n\Omega_{es})E_{zs}}{nk_{\parallel rs}v_{tes}}\right|^2\right)\right.\\&\left.+[I_n(\lambda_{es})-I_{n-1}(\lambda_{es})]\left(2\lambda_{es}|E_{ys}|^2-n\left|E_{xs}+iE_{ys}+\frac{\sqrt{\lambda_{es}}(\omega-n\Omega_{es})E_{zs}}{nk_{\parallel rs}v_{tes}}\right|^2\right)\right].\end{aligned}\tag{3}$$

Here, $\omega_{pes}^2\equiv q_i^2n_s/\epsilon_0m_e$ is the square of the electron plasma frequency, with ϵ_0 the permittivity of free space and q_i the elementary charge. The argument of the modified Bessel function $I_n(\lambda_{es})$ is $\lambda_{es}\equiv k_{\perp rs}^2v_{tes}^2/\Omega_{es}^2$, and $\xi_{ns}\equiv(\omega-k_{\parallel rs}V_{\parallel es}-n\Omega_{es})/\sqrt{2}k_{\parallel rs}v_{tes}$ is the real part of the conventional argument of the plasma dispersion function [36]. The quantity $v_{tes}^2\equiv T_{es}/m_e$ is the square of the electron thermal velocity, $\Omega_e\equiv q_e|\mathbf{B}_{0s}|/m_e$ is the electron cyclotron frequency, and $k_{\parallel rs},k_{\perp rs}$, and $V_{\parallel es}$ are the components of $\mathbf{k}_{rs}$ and $\mathbf{V}_{es}$ parallel or perpendicular to the background magnetic field $\mathbf{B}_{0s}$. The RF field amplitudes $\{E_{xs},E_{ys},E_{zs}\}$, obtained from GENRAY along discrete ray trajectories, are expressed in a coordinate system whose basis vectors are constructed from $\mathbf{k}_{rs}$ and $\mathbf{B}_{0s}$ (see Ref. [21]). Quantities marked with the s subscript may vary on the spatial and temporal scales of extended MHD and are not generally constant along ray trajectories.[1] Also, because of the Gaussian dependence of H_e on ξ_{ns}, it is typical that every term but one in the infinite sum is approximately zero everywhere. Furthermore, the remaining term is typically negligible except at those points along the RF trajectories where $\xi_{ns}\approx0$. This behavior is a physical consequence of the Doppler-shifted RF wave frequency locally aligning with a cyclotron resonance; relatively narrow spatial localization of momentum and energy sources arising from the RF ensues.

In NIMROD, Eq. (1) is implemented by (a) summing over species to yield an expression for the total momentum, and (b) approximating the electron equation to form a generalized Ohm's law. For the total momentum, we obtain

$$m_in_s\left(\frac{\partial\mathbf{V}_s}{\partial t}+(\mathbf{V}_s\cdot\nabla)\mathbf{V}_s\right)=\mathbf{J}_s\times\mathbf{B}_{0s}-\nabla p-\nabla\cdot\boldsymbol{\Pi}+\mathbf{k}_{rs}H_e,\tag{4}$$

where $\mathbf{V}_s\approx\mathbf{V}_{is}$ is the center-of-mass flow velocity, $p\approx n_s(T_i+T_e)$ is the generalized pressure, $\mathbf{J}_s\equiv q_in_s(\mathbf{V}_{is}-\mathbf{V}_{es})$ is the plasma current, and $\boldsymbol{\Pi}\approx\boldsymbol{\Pi}_i+\boldsymbol{\Pi}_e$ is the generalized stress tensor. For Ohm's law, various approximations to

$$\mathbf{E}_{0s}=-\mathbf{V}_s\times\mathbf{B}_{0s}+\eta\mathbf{J}_s+\frac{\mathbf{J}_s\times\mathbf{B}_{0s}}{n_sq_i}+\frac{\mathbf{k}_{rs}H_e}{n_sq_i}+\frac{m_e}{n_sq_i^2}\left[\frac{\partial\mathbf{J}_s}{\partial t}+\nabla\cdot[\mathbf{J}_s\mathbf{V}_s+\mathbf{V}_s\mathbf{J}_s]-\frac{q_i}{m_e}[\nabla(n_sT_e)+\nabla\cdot\boldsymbol{\Pi}_e]\right]\tag{5}$$

may be employed within NIMROD to include resistive dissipation (with resistivity η), the dynamics of Hall MHD, inertia and two-fluid advection, and the neoclassical physics contained in $\boldsymbol{\Pi}_e$. A good overview of the physics inherent in these approximations is provided in Ref. [37]. Much of this physics introduces computational difficulties [38] which are unrelated to the physics of TM feedback stabilization; we are interested primarily in a numerical implementation of the quasilinear terms in Eqs. (2), (4), and (5) and this can be computed in the resistive MHD limit. We will thus use model closures (as in Ref. [39]) with kinematic viscosity ν and anisotropic heat diffusivities $\kappa_\parallel,\kappa_\perp$, together with a simplified Ohm's law,

$$\boldsymbol{\Pi}=-m_in_s\nu\nabla\mathbf{V}_s,\tag{6}$$

$$\mathbf{q}_\alpha=-\frac{3n_s}{2}[\kappa_\parallel\hat{b}\hat{b}+\kappa_\perp(\mathbf{I}-\hat{b}\hat{b})]\cdot\nabla T_\alpha,\tag{7}$$

[1] Note that although we have previously described the "short-wavelength, rapid variation" of the RF fields, the RF field *amplitudes* $\{E_{xs},E_{ys},E_{zs}\}$ of Eq. (3) vary only on the longer spatiotemporal scales of the MHD quantities; the short-wavelength, rapid variation arises from a complex *phase* function $\exp(i\mathbf{k}_s\cdot\mathbf{x}-\omega t)$ which is multiplied by these amplitudes to produce the total RF field. Ref. [21] can be consulted for more details.

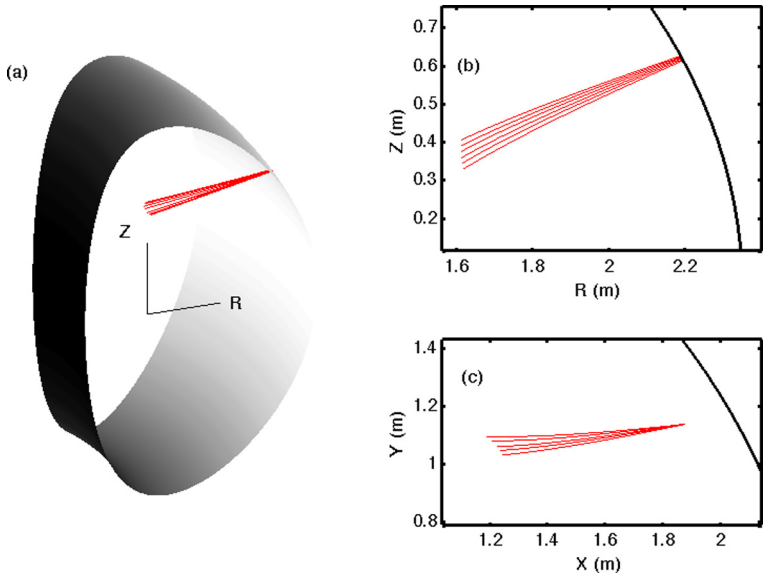

Fig. 1. Typical trajectories of a few RF rays in a bundle are shown (a) in a three-dimensional cutaway section of the torus; (b) as projected into the poloidal (R–Z) plane of the torus; and (c) as projected into the midplane ($Z = 0$, i.e. the view from the top) of the torus. For frequencies near electron cyclotron resonance, trajectories are typically nearly straight except near the constant-R resonant surface of the tokamak where power is deposited.

$$\mathbf{E}_{0s} = -\mathbf{V}_s \times \mathbf{B}_{0s} + \eta \mathbf{J}_s + \frac{\mathbf{k}_{rs} H_e}{n_s q_i}, \tag{8}$$

wherein $\hat{b} \equiv \mathbf{B}_{0s}/|\mathbf{B}_{0s}|$ is the direction of the background magnetic field and $\mathbf{I}$ is the unit tensor. Numerical values for $\nu, \kappa_\parallel, \kappa_\perp$, and η, along with other relevant numerical parameters, will be given in Appendix A.

In NIMROD, plasma densities, temperatures, and flows, together with the associated electromagnetic fields, are spatially represented as continuous quantities. Two-dimensional, high-order finite elements form the basis for these fields in the poloidal plane, while a discrete Fourier representation is used to represent their toroidal variation. The quasilinear source terms of Eqs. (2), (4), and (8) must also be represented in this basis.

2.2. RF equations and their GENRAY implementation

The RF equations we use in this work were also derived formally in Ref. [21], and are valid for the high frequencies and small length scales associated with eikonal-approximated wave propagation through complex media. Their solutions are considerably different from the continuous fields which arise from solutions of the extended MHD equations. The RF wave equation can be written as a matrix operation on the rapidly varying RF electric field vector:

$$\left[\mathbf{N}_s \mathbf{N}_s + \mathbf{I}(1 - N_s^2) + \sum_\alpha \chi_\alpha \right] \cdot \mathbf{E}_s \equiv \mathbb{D}_s \cdot \mathbf{E}_s = 0, \tag{9}$$

wherein $\mathbf{N}_s \equiv \mathbf{k}_s c/\omega$ is the complex refractive index of the plasma (with $\mathbf{k}_s$ the complex wavevector and c the speed of light), and χ_α is a species susceptibility tensor whose hermitian and antihermitian components vary in time and space and satisfy the relation $\chi_\alpha = \chi_\alpha^h + i\chi_\alpha^a$. The precise form of χ_α is derived in Ref. [21], which also discusses in detail how the quantity H_e is obtained from the antihermitian component χ_α^a. Equation (9) is solved by tracking surfaces of constant phase $\psi = \mathbf{k}_s \cdot \mathbf{x} - \omega t$ along their characteristic trajectories. Calculating the determinant $\det[\mathbb{D}_s] \equiv D_s(\mathbf{x}, t, \omega, \mathbf{k}_s)$ and expanding assuming that imaginary components of $\mathbf{k}_s$ and D_s are small, one may derive the ray tracing equations [20] in the form

$$\frac{d\mathbf{k}_{rs}}{dt} = \frac{\partial D_{rs}}{\partial \mathbf{x}} \Big/ \frac{\partial D_{rs}}{\partial \omega}, \tag{10}$$

$$\frac{d\mathbf{x}}{dt} = -\frac{\partial D_{rs}}{\partial \mathbf{k}_{rs}} \Big/ \frac{\partial D_{rs}}{\partial \omega}, \tag{11}$$

wherein $D_{rs} = \mathrm{Re}[D_s(\mathbf{x}, t, \omega, \mathbf{k}_{rs})]$ and the t-dependence reflects the slow (relative to the RF period) alteration of characteristic ray trajectories, and the linear dispersion thereon, as the plasma through which the RF passes evolves on the slow MHD timescales. More general ray tracing techniques [40] may also be employed if the imaginary terms noted above are not small, but the ensuing solutions are still the trajectories of surfaces of constant phase. A typical trajectory plot is shown in Fig. 1, together with various projections of this path into the poloidal and x–y planes. Electric field component ampli-

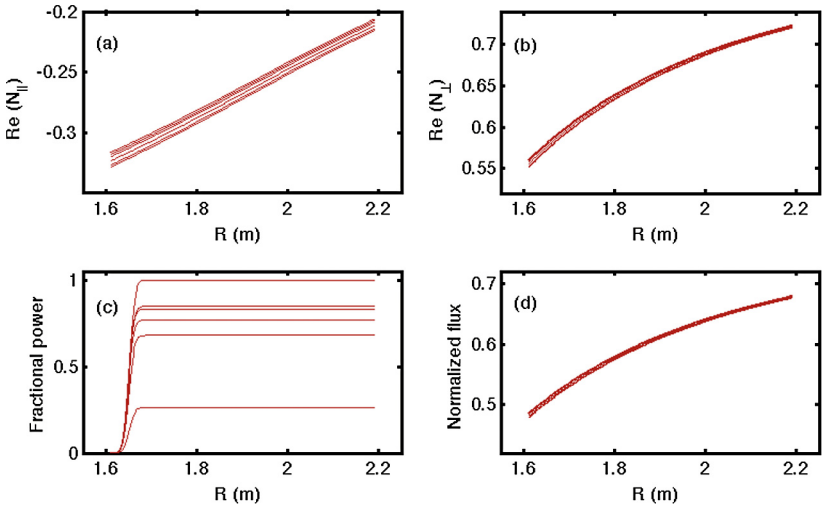

Fig. 2. Variation with respect to major radius, for a subset of rays in the ray bundle, of (a) the real part of the parallel index of refraction $N_{\|}$ $(=k_{\|rs}c/\omega)$; (b) the real part of the perpendicular index of refraction $N_{\perp}$ $(=k_{\perp rs}c/\omega)$; (c) the fractional power content P/P_{max}, where P_{max} is the power of the central ray in the bundle as it exits the RF launcher, and (d) the normalized electromagnetic flux. Rays propagate from larger to smaller R values; the deposition of power in the plasma is indicated by the sharp drop-offs in (c) as the rays intersect the resonant surface.

tudes, local power content, components of the group velocity ($\equiv \mathbf{v}_g$), and other physical variables can also be determined pointwise along these trajectories. GENRAY also calculates the normalized flux

$$\Gamma = \frac{|\mathbf{\Gamma}_P|}{\epsilon_0 c|\mathbf{E}_0|^2} \tag{12}$$

in terms of normalized (complex) RF electric fields $\{e_{xs} \equiv E_{xs}/|\mathbf{E}_0|,\ e_{ys} \equiv E_{ys}/|\mathbf{E}_0|,\ e_{zs} \equiv E_{zs}/|\mathbf{E}_0|\}$; here,

$$\mathbf{\Gamma}_P = \left(\frac{|\mathbf{B}_s|^2}{2\mu_0} + \frac{\epsilon_0 |\mathbf{E}_s|^2}{2} + \sum_\alpha \frac{\epsilon_0}{2} \mathbf{E}_s^* \cdot \frac{\partial}{\partial \omega}(\omega \chi_\alpha^h) \cdot \mathbf{E}_s \right) \mathbf{v}_g. \tag{13}$$

The square of the normalizing field amplitude $|\mathbf{E}_0|$ is proportional to the local power content P of the wave,

$$P \propto \Gamma \epsilon_0 c |\mathbf{E}_0|^2. \tag{14}$$

GENRAY also calculates P along the trajectories. This quantity will be used later to match the input power of the RF source (i.e. the gyrotron power entering the plasma chamber) to the power deposited into the plasma.

For frequencies near electron cyclotron resonance, a conceptual picture of the salient RF physics can often be attained by tracing only a few rays in accordance with Eqs. (10)–(11). For these frequencies, the ray trajectories are generally straight except near the relevant resonant layer; near this layer, RF power is deposited and the rays are reflected. Trajectory, power deposition, and flux plots akin to those in Fig. 2 convey the general behavior of the injected RF fields. Such datasets, however, represent only discrete *samples* of a continuous solution – although the physical RF fields are defined everywhere, tracing a few rays provides values for the continuous physical solution only along the few corresponding trajectory paths. Approximate continuous solutions are obtained only by taking many samples (in this case, by gathering data along many rays) and then considering the global properties of the ensuing large dataset. In particular, the localized deposition of RF power in the plasma – associated with changes in the local power content along the many sampled rays – is a collective property of the set of *all* of the rays. We will call this set of rays the "ray bundle"; our goal is to assess the properties of the bundle (e.g. number/spacing of rays) which are needed for accurate numerical calculations.

To illustrate, suppose that $p^2 = N$ rays (for integer p) are used in GENRAY to model vacuum plane wave propagation in the $\hat{z}$ direction. Assuming periodic boundary conditions and characteristic width a in the x and y directions, we place the wave launcher in the $z = 0$ plane pointing toward positive z and concern ourselves only with the intervals $x \in [0, a]$, $y \in [0, a]$, and $\{z \geq 0\}$. Let the rays be spaced such that the initial position of the i-th such ray is given by

$$(x_i, y_i) = \left(\frac{ai}{p} - \frac{a}{2p} - a\left\lfloor \frac{i}{p} - \frac{1}{2p} \right\rfloor, \frac{a}{p}\left\lceil \frac{i}{p} \right\rceil - \frac{a}{2p} \right) \tag{15}$$

where $\lfloor \ldots \rfloor$ and $\lceil \ldots \rceil$ are the floor and ceiling functions respectively (this method of spacing is illustrated in Fig. 3). The dispersion relation $D_s(\mathbf{x}, t) = k_z^2 c^2 - \omega^2 = 0$ yields, from the ray tracing equations (10)–(11), the results $\partial k_z/\partial t = 0$; $\partial z/\partial t = k_z c^2/\omega = c$; the plane waves propagate in the $\hat{z}$ direction at the speed of light and with constant wavelength. Without loss of generality, we assume the fields take the form $\mathbf{E} = E_{0s} \exp(i\omega[z - ct]/c)\hat{x}$, $c\mathbf{B} = E_{0s} \exp(i\omega[z - ct]/c)\hat{y}$ (i.e.,

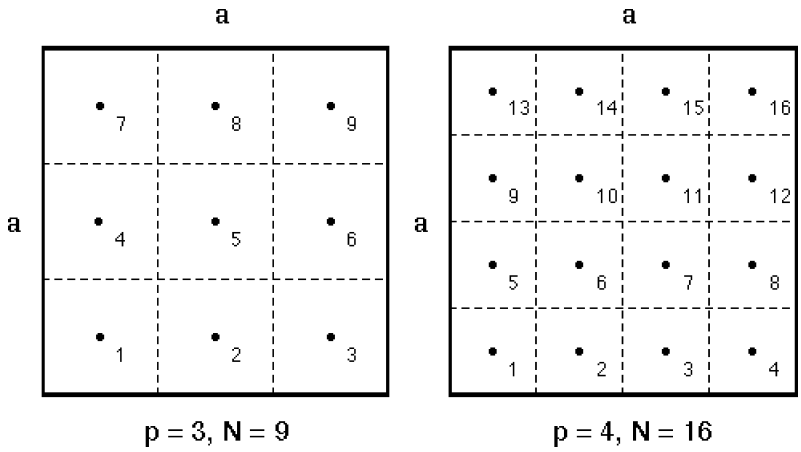

Fig. 3. Spacing and labeling of the $p^2 = N$ initial points in the x–y plane for various values of N, in the plane wave example. Waves propagate in the z-direction (out of the page); the spacing of the rays is chosen such that the cross-sectional area-perpendicular-to-propagation is equal for each ray. Consequently, while the fraction of total power P_0 assigned to each ray ($= P_0/N$) varies with N, the power flux Γ_i through this cross-sectional area remains constant ($= P_0/a^2$) for any value of N at all $z > 0$.

infinitely slowly varying amplitude functions which multiply rapidly varying phase functions). Because the plane wave carries a time-averaged total power P_0, the local power content along any individual ray trajectory is thus (in the absence of dissipation) a constant fraction P_0/N of this total power. Though this local power content is not a physically relevant quantity (since it varies with N), a local power *flux* which is independent of N can be determined along individual ray trajectories. It is evident from Fig. 3 that each ray is effectively responsible for modeling the physics within a box semi-infinite in z and with cross-sectional area

$$A = \left(\frac{a}{p}\right)^2 = \frac{a^2}{N} \tag{16}$$

within an arbitrary ($z > 0$) x–y plane. Consequently, the local power flux Γ_i along the i-th ray trajectory can be directly calculated by taking the ratio of the local power content to the (constant) cross-sectional area-perpendicular-to-propagation associated with that trajectory. We obtain

$$\mathbf{\Gamma}_i = \frac{P_0/N}{a^2/N}\hat{z} = \frac{P_0}{a^2}\hat{z}\,, \tag{17}$$

which correctly represents the total power flux along the trajectory regardless of the number of rays involved in the calculation. Further, because this total power flux is constant, we may also write

$$\mathbf{\Gamma}_P = \epsilon_0 c E_{0s}^2 \hat{z} = \frac{P_0}{a^2}\hat{z} = \frac{P_0/N}{a^2/N}\hat{z} \tag{18}$$

to reveal that the factor relating the two sides of Eq. (14), for a single trajectory, is the cross-sectional area a^2/N associated with that trajectory at every point.

What if the N rays of this example were not regularly spaced or did not contain equal shares of the total power P_0? In either case, we know that the power flux $\mathbf{\Gamma}_i$ remains constant since the nature of the plane wave we are modeling is unchanged. It follows that the local power content associated with a particular ray is inextricably linked with an effective area – i.e., some portion of the cross-section of the ray bundle – which correctly relates the local power in the ray with its contribution to the local flux. In more general cases, this flux may be an unknown, spatially varying quantity, but the use of an adequately large number of densely packed rays (such that the average ray separation is much less than the scale length on which the variation occurs) allows the relationship between local power content and local flux to persist, so long as the effective area can be determined. In physically realistic cases, the rays in question are nonparallel and their local power content varies along trajectories as power is dissipatively transferred to the plasma. In such cases, the local power flux and effective area may likewise vary along the trajectories. We will demonstrate, in Section 3, that the effective area can be calculated by applying methods of computational geometry to the collective data of the ray bundle.

In summary, GENRAY solutions to the RF ray tracing equations are obtained only along discrete spatial trajectories which are the characteristics of ordinary differential equations. Variation in the plasma wavenumber, electric field components, fluxes, and local power content along the trajectories is also computed. The collective geometric properties of the discrete trajectory paths must be used to map this discrete data onto NIMROD's continuous representation. The use of a sufficiently large number (typically hundreds) of rays enables the detailed spatial resolution needed to construct a smooth model for the RF-induced effects in NIMROD's fluid evolution equations.

3. Computational geometry

We have discussed how the coupled RF/MHD system can be modeled using NIMROD to evolve the fluid variables and GENRAY to calculate the RF physics. Because RF rays propagate through the MHD profiles, alterations to their trajectories (and the corresponding physical variables sampled along these trajectories) will ensue as the plasma evolves. These variations, though nonnegligible, are generally minor, and the transfer of updated profile data to be used in new trajectory calculations (i.e. the transfer of MHD profile data to GENRAY) is straightforward. The opposite process – namely, the transfer of RF data to NIMROD – requires us to map the discrete ray data to NIMROD's continuous representation by construction of the effective area elements. Brief, heuristic explanations of the computational geometry methods we will use are provided in this section.

We again consider a bundle of N nearly parallel rays emitting from a point representing the RF launcher; the trajectories through the plasma are calculated from Eqs. (10)–(11). We represent the trajectory of the pth ray by $\mathbf{x}_p(s)$, where s parametrizes the distance along the trajectory and is discretized such that $s \sim s_m = m\Delta s$ for integer m. At every point $\mathbf{x}_p(s_m)$ at which the physics of RF deposition is nonnegligible [a relatively small number of points due to the $\exp(-\xi_{ns}^2)$ factor in Eq. (3)], we determine a value for the area perpendicular to the direction of propagation as follows:

1. Determine the local direction of ray propagation $\hat{\mathbf{n}}_{p,m}$ by interpolation,

$$\hat{\mathbf{n}}_{p,m} = \frac{\mathbf{x}_p(s_{m+1}) - \mathbf{x}_p(s_{m-1})}{|\mathbf{x}_p(s_{m+1}) - \mathbf{x}_p(s_{m-1})|}. \tag{19}$$

2. Define the plane $P_{p,m}$ which contains the point $\mathbf{x}_p(s_m)$ and is perpendicular to $\hat{\mathbf{n}}_{p,m}$. Some portion of this plane will be the desired area.
3. Construct the set $S_{p,m}$ of points in $P_{p,m}$ which are intersected by rays in the bundle. This set includes $\mathbf{x}_p(s_m)$ and may include as many as $N-1$ other points (though there may be fewer since an arbitrary ray does not necessarily intersect an arbitrary plane). Fundamentally, $S_{p,m}$ is a set of coplanar points akin to the dots in Fig. 4(a), and we are assigning an area to one such point by considering the position of its neighbors.
4. Construct the Voronoi tessellation of $S_{p,m}$ (the solid black lines in Fig. 4(a)). The Voronoi tessellation decomposes the plane $P_{p,m}$ into separate regions corresponding to the individual members of $S_{p,m}$. All points within a region are closer to their particular member of $S_{p,m}$ than they are to any other member of the set. This defines 'neighborhoods' about the members of $S_{p,m}$ to which the points in the plane most naturally belong.
5. Calculate the area $A_{p,m}$ of the region of the Voronoi tessellation corresponding to the desired point.

For a certain subset of points in $S_{p,m}$ this method fails to yield finite areas. An arbitrary point in the plane $P_{p,m}$ may be infinitely far away from the ray bundle and yet be nearer to one point in the bundle than to any others; consequently, regions of the Voronoi tessellation which correspond to the outermost points of $S_{p,m}$ may extend to infinity. To circumvent this issue, ghost boundary points can be appended to $S_{p,m}$ in the following manner:

3(a). Determine the convex hull of $S_{p,m}$. The convex hull is effectively the outer boundary of $S_{p,m}$ in the plane $P_{p,m}$. (Heuristically, envision surrounding $S_{p,m}$ with a rubber band and then allowing the band to contract until it is taut.) The vertices of the convex hull are the points where the previous method fails, as the semi-infinite regions of Fig. 4(a) suggest.

3(b). Construct the Delaunay triangulation of $S_{p,m}$, the dashed lines of Fig. 4(b). The Delaunay triangulation subdivides the region of $P_{p,m}$ within the convex hull into triangles whose minimum interior angles are collectively maximized. Its vertices form a dual mesh complementary to the vertices of the Voronoi tessellation. A circle which circumscribes the vertices of an arbitrary triangle in the Delaunay triangulation will contain no other points in $S_{p,m}$; long, skinny triangles are thus discouraged.

3(c). Each line segment between members of $S_{p,m}$ on the convex hull forms an edge of a triangle in the Delaunay triangulation. For each such segment, mirror the corresponding triangle across the segment in the plane (effectively 'folding' the triangle across the convex hull, as in Fig. 4(c)). This creates a new triangle, two of whose vertices are already in the set $S_{p,m}$. We add the third vertex point (marked with the + symbol in the figure) to a new set $S'_{p,m}$.

3(d). When all items of $S'_{p,m}$ have been determined, append $S'_{p,m}$ to $S_{p,m}$ and proceed to step 4 above. Vertices of the convex hull of the expanded set will all be members of $S'_{p,m}$, and thus, only these 'ghost boundary points' will correspond to unbounded regions of the Voronoi tessellation (see Fig. 4(d)). All members of the original set $S_{p,m}$ can then be put into correspondence with bounded regions of this tessellation, and the area of these regions can be calculated.

The computational geometry calculations (Delaunay triangulation, convex hull identification, Voronoi tessellation) are carried out using the QHULL software package [41]. Using the foregoing methods, representations for the source terms $\mathbf{k}_{rs}H_e$, $\mathbf{k}_{rs}H_e/n_sq_i$, and ωH_e can be obtained from the discrete data computed along the ray trajectories. We explicitly

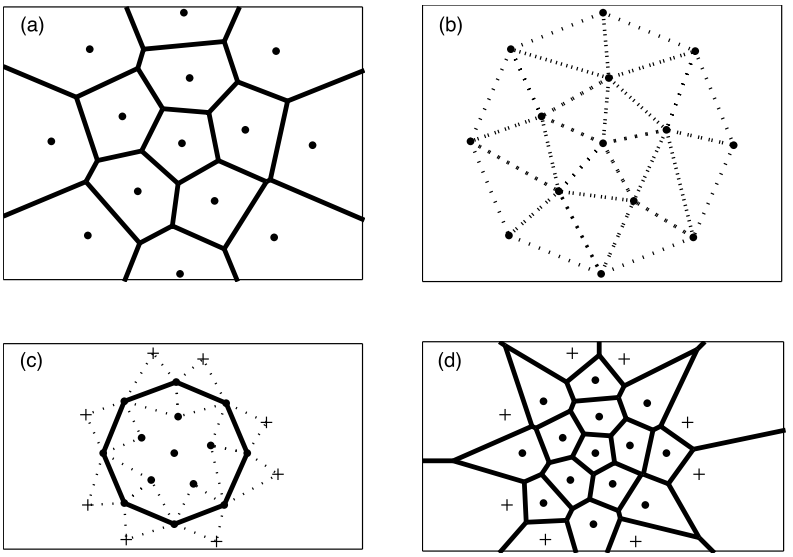

Fig. 4. (a) A set of coplanar points representing the intersection of a ray bundle with a plane are shown together with the Voronoi tessellation (solid black lines) corresponding to these points. Note that the Voronoi regions corresponding to the outermost points extend to infinity. (b) Delaunay triangulation of the points in (a). (c) Convex hull (solid black lines) of the points in (a), together with the reflection across the hull of the Delaunay triangles whose edges lie on the convex hull. Ghost points (+ symbols) are thus created outside the convex hull. (d) Voronoi tessellation of the set of original and ghost points. Note that all original data points now correspond to Voronoi regions of finite area.

write H_e in terms of variables known along these trajectories (here parametrized by the discrete quantity s_j) in the form

$$
\begin{aligned}
H_e(s_j) = \sum_{n=-\infty}^{\infty} & \frac{q_i^2 n_s(s_j)\sqrt{\pi}}{4\omega^2\sqrt{m_e}} \frac{[c - N_{\parallel rs}(s_j)V_{\parallel es}(s_j)]}{\sqrt{2}N_{\parallel rs}(s_j)\sqrt{T_{es}(s_j)}} \frac{P(s_j)}{\Gamma(s_j)\epsilon_0 c A(s_j)} \\
& \times \exp\left(-\frac{[\omega c m_e - \omega N_{\parallel rs}(s_j)V_{\parallel es}(s_j)m_e + n q_i |\mathbf{B}_{0s}(s_j)|c]^2}{2\omega^2 N_{\parallel rs}(s_j)^2 T_{es}(s_j) m_e}\right) \exp\left(-\frac{N_{\perp rs}(s_j)^2\omega^2 T_{es}(s_j)m_e}{q_i^2|\mathbf{B}_{0s}(s_j)|^2c^2}\right) \\
& \times \left\{\left[I_n\left(\frac{N_{\perp rs}(s_j)^2\omega^2 T_{es}(s_j)m_e}{q_i^2|\mathbf{B}_{0s}(s_j)|^2c^2}\right) - I_{n+1}\left(\frac{N_{\perp rs}(s_j)^2\omega^2 T_{es}(s_j)m_e}{q_i^2|\mathbf{B}_{0s}(s_j)|^2c^2}\right)\right]\right. \\
& \times \left(\frac{-2\omega N_{\perp rs}(s_j)\sqrt{m_e}\sqrt{T_{es}(s_j)}}{c q_i |\mathbf{B}_{0s}(s_j)|}|e_{ys}(s_j)|^2\right. \\
& \left. + n\left|e_{xs}(s_j) - i e_{ys}(s_j) - \frac{N_{\perp rs}(s_j)[m_e\omega + n q_i|\mathbf{B}_{0s}(s_j)|]e_{zs}(s_j)}{n N_{\parallel rs}(s_j) q_i |\mathbf{B}_{0s}(s_j)|}\right|^2\right) \\
& + \left[I_n\left(\frac{N_{\perp rs}(s_j)^2\omega^2 T_{es}(s_j)m_e}{q_i^2|\mathbf{B}_{0s}(s_j)|^2c^2}\right) - I_{n-1}\left(\frac{N_{\perp rs}(s_j)^2\omega^2 T_{es}(s_j)m_e}{q_i^2|\mathbf{B}_{0s}(s_j)|^2c^2}\right)\right] \\
& \times \left(\frac{-2\omega N_{\perp rs}(s_j)\sqrt{m_e}\sqrt{T_{es}(s_j)}}{c q_i |\mathbf{B}_{0s}(s_j)|}|e_{ys}(s_j)|^2\right. \\
& \left.\left. - n\left|e_{xs}(s_j) + i e_{ys}(s_j) + \frac{N_{\perp rs}(s_j)[m_e\omega + n q_i|\mathbf{B}_{0s}(s_j)|]e_{zs}(s_j)}{n N_{\parallel rs}(s_j) q_i |\mathbf{B}_{0s}(s_j)|}\right|^2\right)\right\},
\end{aligned}
\tag{20}
$$

where $P(s_j)$ is the local power content at discrete points along the trajectory and $A(s_j)$ is the computed area of the corresponding Voronoi tessellation at these points.

Having obtained a discretized representation for the source terms along the GENRAY trajectories, we turn to a discussion of how these source terms are to be transformed to NIMROD's spectral/finite element representation.

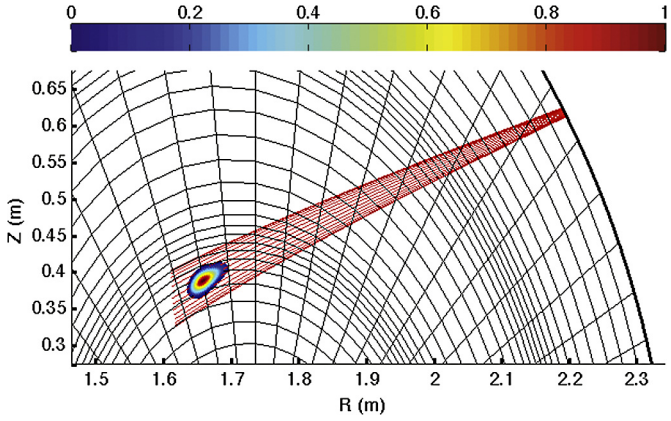

Fig. 5. The normalized function H_e projected into the poloidal plane, akin to plot (b) of Fig. 1, together with corresponding projections of the ray trajectories. The grid boundaries corresponding to finite element cells in the poloidal plane are also displayed; within the finite element cells, high-order bivariate polynomials model physical quantities. Grid packing has been employed and results in nonuniform cell sizes. The grid resolution shown (48 radial and 64 poloidal grid cells with fourth-order polynomials) is adequate to capture variation in H_e in the plane.

4. Interpolation to finite element and spectral basis functions

NIMROD uses a discrete spectral representation in the toroidal direction. For a fixed position in the R–Z poloidal plane, the toroidal representation has the form

$$A(R,Z,\phi) \approx A_n(R,Z) = \tilde{A}_0(R,Z) + \sum_{p=1}^{N} [\tilde{A}_p(R,Z)e^{i\pi np/N} + \tilde{A}_p^*(R,Z)e^{-i\pi np/N}], \tag{21}$$

and represents the phases and amplitudes of N modes in the Fourier space. Functionally, this representation is equivalent to a sample of the toroidally varying function taken on $2N$ discrete, regularly spaced poloidal planes in real space [at values $\phi = \pi n/N$, $n \in -N \ldots (N-1)$]. Aliasing errors are prevented in NIMROD by zeroing Fourier components for $|n| > \lfloor 2N/3 \rfloor$; however, this does not fundamentally alter the representation in real space. Within poloidal planes, physical quantities are expanded using high-order finite elements. Grid packing of the finite elements can be used to locally increase resolution in regions of interesting physics within the poloidal plane. In contrast, the discrete spectral representation samples this variation only at uniformly spaced toroidal intervals. Often, a relatively small number of Fourier modes are needed in NIMROD tokamak simulations because of the dominant toroidal magnetic field.

For ECCD injection in a tokamak, however, the $\exp(-\xi_{nr}^2)$ factor of Eq. (3), together with the typically narrow width of the ECCD launch cone [42], lead to power deposition over extremely narrow spatial regions (see Figs. 5 and 7). The momentum/energy/current sources induced in NIMROD will be correspondingly localized. Within the poloidal plane, grid packing – as shown in Fig. 5 – is used, together with biquartic polynomials on a finite element grid (a 48×64 grid is typical for NIMROD runs) to enable resolution of H_e in this plane. However, the localized toroidal variation of this function presents some difficulty. Quantitative estimates for toroidal resolution requirements can be made by considering the function

$$A(\phi) = \begin{cases} \dfrac{1}{\phi_c} \cos\left(\dfrac{\pi(\phi-\phi_0)}{2\phi_c}\right)^2, & |\phi-\phi_0| \le \phi_c \\ 0, & \text{otherwise,} \end{cases} \tag{22}$$

which is a smooth, localized function with unit area and variable toroidal width. This function is shown in Fig. 6a; both the function and its derivative are zero at the bounding points $|\phi - \phi_0| = \phi_c$. Using the Fourier representation for a continuous function on a finite interval, one can write

$$\tilde{A}_l = \frac{1}{2\pi}\int_0^{2\pi} A(\phi)e^{il\phi}d\phi = \frac{e^{-il\phi_0}}{2\pi}\frac{\sin(\pi y_l)}{(1-y_l^2)\pi y_l} \tag{23}$$

where $y_l \equiv l\phi_c/\pi$. These Fourier coefficients (shown in normalized form in Fig. 6(b)) fall off rapidly for large y_l, and never exceed 0.03 of the fundamental ($l = 0$) amplitude for $y_l > 2$. Suppose that after dealiasing, we retain all Fourier components for which $y_l \le 2$ to ensure that the general functional form of $A(\phi)$ is approximately preserved. Demanding that our chosen relation $y_l = l\phi_c/\pi \le 2$ is at least as restrictive as the dealiasing relation $l < 2N/3$, we have

$$\frac{2\pi}{\phi_c} < \frac{2N}{3}; \quad 6\left(\frac{2\pi}{2\phi_c}\right) < 2N. \tag{24}$$

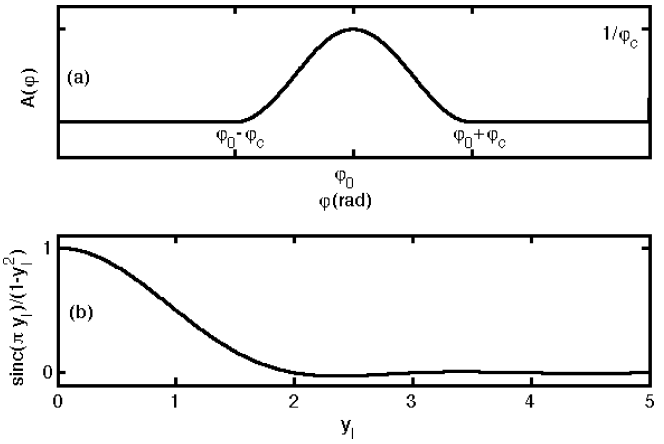

Fig. 6. (a) General function for describing toroidal variation using parameters $\{\phi_c, \phi_0\}$, as in Eq. (22). (b) Fourier coefficients [neglecting a phase factor and $(2\pi)^{-1}$ normalization constant] corresponding to this form. The coefficients are non-negligible for $y_l < 2$; dealiasing below this limit will introduce inaccuracy.

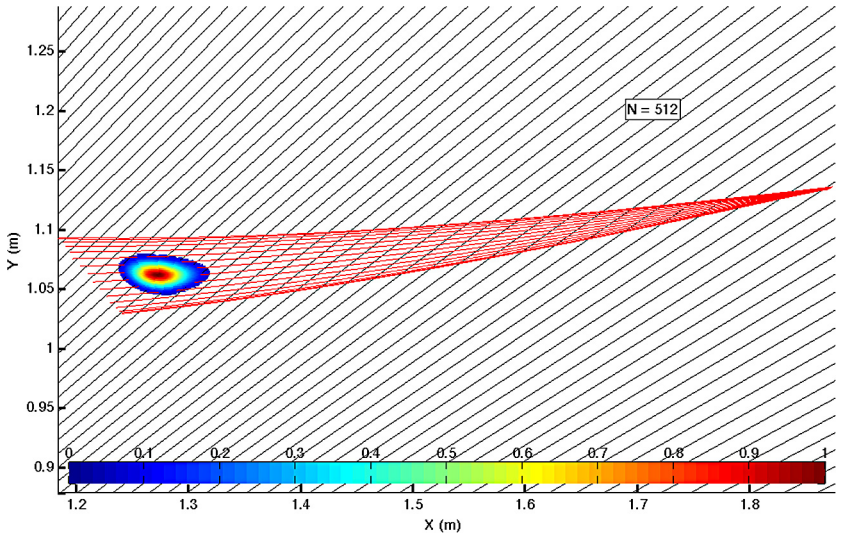

Fig. 7. Midplane projection, akin to plot (c) of Fig. 1, of the normalized function H_e atop the ray trajectories. Planes of constant ϕ corresponding to an $N = 512$ Fourier sampling, which resolves the centimeter-scale toroidal variation of H_e, are also displayed.

Because $2N \equiv 2^m$ is the number of discrete poloidal planes on which the function is sampled in real space, and $2\phi_c/2\pi \equiv s$ is the fractional width of the function relative to the toroidal period ($0 \le s \le 1$), we may also write

$$m > \frac{\ln(6)}{\ln(2)} - \frac{\ln(s)}{\ln(2)}, \tag{25}$$

which gives the power-of-two needed to resolve physics occurring on toroidal scales of fractional width s. For Fig. 7, in which the angular width of the RF deposition is approximately $\pi/64$, this corresponds to $s = 1/128$ and $m > 9.58$. Thus, a minimum of $N = 2^9 = 512$ toroidal Fourier modes (corresponding to 1024 poloidal planes) are needed.

NIMROD has demonstrated scaling capability on high-performance clusters and is fully capable of running with hundreds of toroidal Fourier modes. Weak scaling results obtained on the Cray XK7 (Titan, at ORNL) are shown in Fig. 8 for $N = 2^p$ such modes ($p \in [6 \ldots 9]$) in simulations which make use of 5632 to 43,776 cores. The NIMROD time-advance loop kernel time is shown, along with non-overlapping subset-kernel times from SuperLU (preconditioner), finite-element matrix, and finite-element right-hand-side computations. Good scaling is seen for both the SuperLU and FE matrix kernels, since the latter performs no communication and the SuperLU processor grid size is unchanged between runs. As additional Fourier modes are added, the number of processors participating in the all-to-all calls for FFTs is increased proportionally to the problem size, as reflected in the increased FE RHS kernel times. Despite an eight-fold increase in problem size, the computation time required for 50 steps of NIMROD's main time loop is increased only by $\sim 35\%$, with the additional time and communication requirements for computing increasingly larger fast Fourier transforms contributing substantively to the increase. Though the weak scaling is imperfect, the increased computation time associated with higher numbers of modes

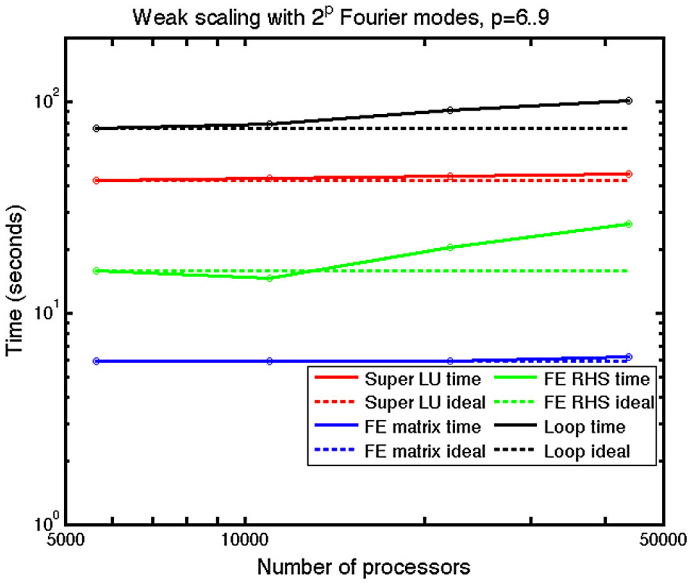

Fig. 8. Weak scaling on the Cray XK7 (Titan) high-performance cluster as the toroidal resolution of the NIMROD simulation is increased commensurate with the processor count, together with ideal scaling metrics. The time required to traverse 50 steps of NIMROD's main time loop is shown, together with non-overlapping subset-kernel times from preconditioning (Super LU), finite-element matrix assembly (FE matrix), and right-hand-side (FE RHS, including the fast Fourier transforms associated with toroidal resolution) computations. With an eight-fold increase in toroidal resolution, total computation time is increased only by ∼35%, deviating from ideal scaling predominantly due to the larger, more communication-intensive Fourier transform operations required at higher resolutions.

is far from prohibitive. It should also be noted that the toroidal resolution required for conventional NIMROD use cases involving simulation of core modes is much lower than was needed for the cases shown in Fig. 8; the high Fourier mode count is required only to toroidally resolve the RF source used in this work.

We turn now to the issue of transforming data associated with discrete rays which intersect a plane into a continuous representation within that plane.

4.1. Poloidal interpolation

As mentioned previously, NIMROD uses a 2D finite-element representation in the poloidal plane. This representation is nodal with coefficients of the 2D basis functions taking on the value of the dependent variables at nodes in the grid. This section contains a description of a method to project GENRAY data onto NIMROD's finite-element nodes. It uses polynomial interpolation along the ray trajectory and an efficient bivariate (2D) interpolation algorithm due to Renka [43,44] in the poloidal plane.

Recall that the spacing of data along a single ray trajectory is dense compared to the density of the rays themselves. Using this fact, we use low-order polynomial interpolation along the rays to project data onto constant-ϕ poloidal planes. This results in relatively sparse, non-uniformly distributed planar data sets. Fig. 9(a) shows the NIMROD finite element node points along with a set of $N_r = 331$ such crossing points (R_i, Z_i) (where $i = 1, \ldots, N_r$) for the underlying function H_e in Eq. (20). Fig. 9(b) contains the same data, but with additional rays added at the periphery of the ray cone ($N_r = 661$) for reasons that will be explained later in this section. A modified Shepard's algorithm [45] is applied to these 2D data sets according to the interpolatory formula:

$$H_e(R, Z) = \sum_{k=1}^{N_r} \bar{W}_k(R, Z) Q_k(R, Z). \tag{26}$$

Here the bivariate, quadratic interpolation function satisfies $Q_k(R_k, Z_k) = H_e(R_k, Z_k)$, and fits values of H_e on a subset of nearest neighbor nodes in a weighted least squares sense. The weight functions satisfy $\bar{W}_k(R_j, Z_j) = \delta_{jk}$ and are inversely proportional to the squared distance between (R, Z) and (R_k, Z_k). The CSHEP2D [44] routine, which implements this interpolation, includes an efficient search algorithm for finding nearest neighbors.

Figs. 9(c) and 9(d) show contours of H_e evaluated using NIMROD's 2D finite-element representation following bivariate interpolation of the data in Figs. 9(a) and 9(b) to grid nodes. The oscillations at the periphery of the ray bundle in Fig. 9(c) are tied to the fact that a limited number of data points were associated with the outer region of the ray cone where H_e is close to zero. For this case, the local interpolatory functions, Q_k, rely too much on data near the center of the ray cone – that is, the subset of nearest neighbors does not well-represent the small values of the function near the exterior of the ray cone. For the contours in Fig. 9(d), additional rays were added in the outer region. Though these rays carry virtually

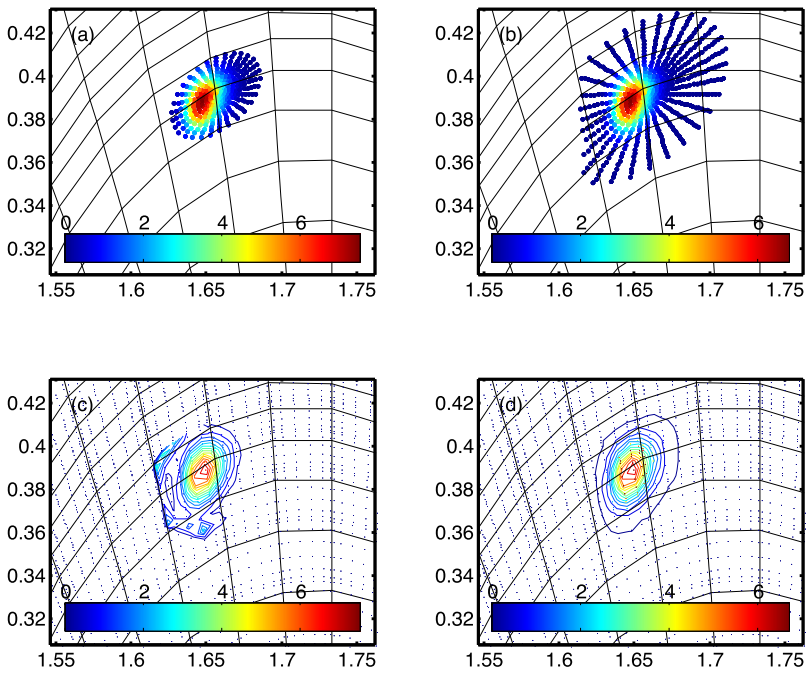

Fig. 9. (a) Value of the function H_e at $N_r = 331$ points where the RF rays intersect a poloidal (constant-ϕ) plane. (b) The same data, but with additional rays added at the periphery of the ray bundle; $N_r = 661$ points. (c) Contours of the bivariate Shepard interpolation of the data in (a) onto NIMROD's Gaussian quadrature points (also shown, together with the finite element grid nodes). Inadequate data at the periphery of the ray cone prevents accurate representation of the vanishing of the H_e function in this region; the Shepard interpolation thus introduces spurious oscillations away from the region of RF deposition. (d) When the interpolation procedure used in (c) is applied to the larger bundle in (b), the additional rays quell the spurious oscillations and yield a smooth, accurate representation of the RF source in the NIMROD basis. Differences in the maximum H_e values between (a/c) and (b/d) arise from differing fractional RF power content within the rays. Wider ray cones compute slightly larger values for the total RF power (by integrating over a wider portion of its angular distribution); thus, rays at a fixed angle will have slightly lower fractional power content for the wider cone.

no power, the near-zero data they contain effectively constrains the Shepard interpolation at the periphery of the ray cone, thus preventing the spurious oscillations.

5. Results and discussion

With accurate models for the RF sources, we are now in position to model the suppression of tearing modes by RF. In this work, as in Ref. [39], we focus on low-β plasmas wherein tearing modes are dominated by resistive (rather than neoclassical) effects; this choice allows us to demonstrate the numerical methods we have developed here while still allowing the use of reasonably simple closures for the extended MHD equations. (Expressions for the closures are given in Appendix A.) In Fig. 10, Fourier components of the magnetic energy (corresponding roughly to the fourth power of island width) are shown in a regime unstable to a magnetic island of helicity $(2, 1)$. Following a period of linear growth, RF waves resonant with the second electron cyclotron (EC) harmonic at the island center are injected, wholly arresting the mode growth. As in Ref. [39], this stabilization arises from modifications of the underlying equilibrium current profile, with consequent modification of the tearing stability parameter Δ', by the injected RF; when the injected RF is shut off, mode growth resumes. As we have previously noted, experimental mitigation and control of tearing modes is generally carried out via sophisticated control systems [12–16]. The approaches developed in this work, in tandem with the theoretical foundations laid in Refs. [17–19, 21], facilitate the development of numerical analogues for such control systems [46]. We anticipate exploring this topic in greater depth in subsequent publications; investigation of the relationship of experimentally observed Fisch and Boozer [47] and Ohkawa [48] currents to RF-induced modifications of the closures is also a topic of interest. The recent development and implementation of continuum closure capabilities in NIMROD [49], via the solution of species drift-kinetic equations, is expected to enable significant progress in these efforts. In particular, the accurate inclusion of critical neoclassical effects such as bootstrap current, via the computation of parallel electron stresses and collisional friction in higher-β plasmas, will be crucial for modeling experimentally relevant NTM dynamics and stabilization in tandem with the methods we have proposed here.

Computational requirements associated with GENRAY and the geometric calculations discussed in Section 3 have not been discussed extensively in this work. For the coupled case presented here, these requirements are negligibly small in comparison with the NIMROD computational requirements. The latter require hundreds to thousands of processors for order-tens of hours to simulate tearing mode growth, while the geometric and ray tracing computations we have presented typically make use of no more than 20–30 processors for order tens of seconds. Increased resolution of the RF sources,

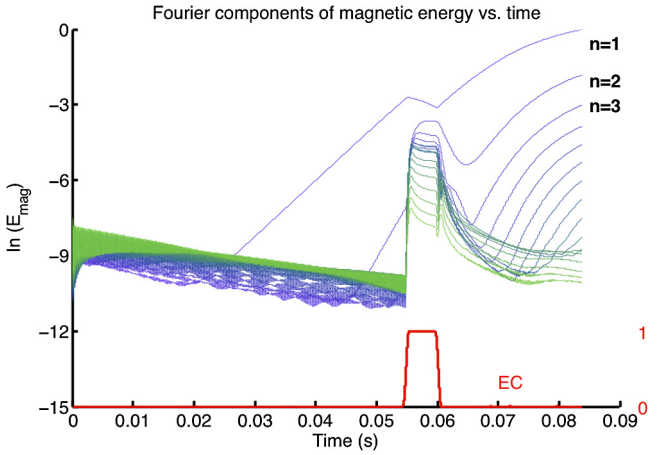

Fig. 10. Fourier components of magnetic energy $B^2/2\mu_0$ are shown during the initial linear growth phase of a resistivity-dominated tearing mode with $(2, 1)$ helicity. The introduction of electron-cyclotron (EC)-resonant RF at the center of the induced magnetic island, using the numerical methods outlined herein, arrests the growth of the mode until the RF is shut off. The development of more sophisticated control algorithms for the RF is a topic of ongoing research.

and thus higher computational requirements, may be necessary in future simulations exploring high-β NTM stabilization scenarios or the interaction of RF with other neoclassical effects. However, the associated need for continuum closure computations in such simulations (as noted in the preceding paragraph), together with the large initial disparity between RF and MHD computation requirements in the coupled case shown here, suggest that computational requirements for such coupled RF/MHD simulations will continue to be dominantly determined by the requirements for the extended MHD code.

In this work, we have demonstrated that physical quantities associated with discrete ray trajectories – the characteristic solutions of ordinary differential equations used in ray-tracing applications – can be interpolated onto continuous finite-element and pseudospectral representations via methods of computational geometry. The use of these methods to compute the effective area-perpendicular-to-propagation associated with any such ODE trajectory, using the collective properties of the ray bundle, forms the basis of the interpolation scheme. The effective areas can then be used in the construction of fluxes and other (not necessarily linear) physical quantities associated with the ray propagation, in the desired continuous representation. Spatial convergence of the interpolation scheme thus relies on the smallness of the inter-ray spacing relative to the scale length of the physics under consideration; we have also demonstrated that the bundle must be broad enough that the interpolation methods do not introduce spurious physics at its edge. Temporal convergence is less of a concern, as the dominant time variation of RF source terms which enter the NIMROD equations arises from alterations to NIMROD profiles on the slower timescales relevant to extended MHD dynamics. For the cases considered here, variation of order 1% in these source terms generally ensues only after order-hundreds of NIMROD timesteps. Simulation of more sophisticated experimental control techniques, such as the use of time-modulated RF or real-time steering of RF launchers, will alter the position and strength of RF sources on more rapid timescales and will thus require better metrics for assessing convergence in the time domain. We are hopeful that such techniques can be further explored in future publications.

Because the ray bundle consists essentially of values attached to an arbitrarily-sized, spatially unstructured mesh which will be interpolated to a spatially structured finite-element mesh, developing quantifiable convergence metrics for spatial resolution of the coupled NIMROD and GENRAY codes is a nontrivial exercise, and no one metric lends itself well to a description of the process. As a general rule, a series of steps must be carried out in order to characterize the scale lengths of the physics to be modeled relative to the various spatial metrics of the two codes. One first focuses on the GENRAY inputs, modifying the spacing of points Δs (using the notation of Section 3) along the ray trajectory until smooth variation of relevant physical quantities (e.g., those in Fig. 2) along this trajectory is obtained. Thereafter, estimates of the parallel scale length $L^{\parallel}$ associated with this smooth variation should be compared with the perpendicular scale length $L^{\perp}_{p,m} \equiv \sqrt{A_{p,m}}$, obtained from the square root of the effective area $A_{p,m}$ corresponding to a given point along the ray trajectory. More rays should be added to the GENRAY computation if the condition $L^{\parallel} \gtrsim L^{\perp}_{p,m}$ is not met. Subsequent interpolation of the converged ray data onto NIMROD's finite-element basis functions generates a finite-element representation for the source term [see Eq. (26)] with its own characteristic scale length L_{source}. Additional rays must be added to the periphery of the GENRAY bundle, and the whole process begun anew, if this source does not smoothly drop to zero outside the deposition region. Finally, comparison of the source scale length with the characteristic spacing of nodes on the NIMROD finite-element grid can be made using plots akin to those of Fig. 9. Spacing between nodes of the bivariate polynomials (in this case, biquartic) associated with individual finite-element grid cells should be sufficiently small that the scale length L_{source} can be well resolved. In general, the process for checking spatial convergence is somewhat cumbersome and does not lend itself well to representation by a single metric. Nevertheless, the foregoing process appears to reliably produce smooth source terms consistent between GENRAY and NIMROD representations.

While we have applied the interpolation techniques discussed herein for the modeling of RF-induced instability mitigation in fusion devices, the methods we have presented can be generally applied to generate new representations for any sufficiently dense bundle of characteristic ODE trajectories. Models of radar, acoustic, optical, or other phenomena satisfying an eikonal approximation may thus also profitably make use of the interpolation methods shown here.

Acknowledgements

This work was supported by the SciDAC Center for Simulation of RF Wave Interactions with Magnetohydrodynamics (SWIM), which jointly received funding from the U.S. Department of Energy under the Fusion Energy Sciences and Advanced Scientific Computing Research programs of the Office of Science, via Cooperative Agreement No. DE-FC02-06ER54899. The authors wish to thank D.B. Batchelor, J.D. Callen, R.W. Harvey, C.C. Hegna, J.R. King (particularly for assistance with the scaling results), S.E. Kruger, J.J. Ramos, D.D. Schnack, D.N. Smithe, and C.R. Sovinec, together with other members of the NIMROD, SWIM, and SciDAC Center for Extended MHD Modeling (CEMM) teams, for support and useful discussions. We also acknowledge, with appreciation, the constructive and helpful feedback provided by the reviewers.

This research used resources of the Oak Ridge Leadership Computing Facility (at Oak Ridge National Laboratory) and the National Energy Research Scientific Computing Center (at Lawrence Berkeley National Laboratory), both of which are supported by the Office of Science of the U.S. Department of Energy under contracts DE-AC05-00OR22725 and DE-AC02-05CH11231 respectively.

Appendix A. Numerical parameters

Exact solutions to the fluid closure problem require the determination of values for heat fluxes ($\mathbf{q}$) and stresses ($\mathbf{\Pi}$) as functions of lower-order fluid moments. In addition to the complexities of the standard closure problem [50], the calculation of a physically consistent closure for the coupled ECCD/MHD model presented here and in Ref. [21] must also take into account the presence of the quasilinear terms in the kinetic equation [Eq. (31) of the latter work]. As previously noted, our focus here is on the numerical implementation of RF-induced momentum and energy sources, and it thus suffices to use simple models for the closure relations and model tearing modes whose dominant physics arises from resistivity, rather than from neoclassical physics. We use Eqs. (6)–(7) with parameters $\nu, \kappa_\parallel$, and $\kappa_\perp \ll \kappa_\parallel$ (respectively, the kinematic viscosity, parallel heat diffusivity, and perpendicular heat diffusivity) taking the respective values 4.2×10^{-1} m^2/s, 4.2×10^{7} m^2/s, and 4.2×10^{1} m^2/s. The underlying Miller equilibrium, obtained from an inverse Grad–Shafranov solve, has a mild pressure gradient [$p \sim (1-\psi)^4$] atop a gauge pressure of 4005.0 Pa; the magnetic profile has on-axis safety factor $q(0) = 1.6$ and $q_{95} = 6.14$. Other parameters relevant to these simulations include the plasma's Lundquist number $S = 2.4 \times 10^6$, the Alfvén and resistive times (respectively $\tau_A \equiv \sqrt{\mu_0 m_i n_i} R_0/B_0 \approx 4.2 \times 10^{-7}$ s and $\tau_R \equiv a^2 \mu_0/\eta \approx 1.0$ s), the plasma resistivity $\eta/\mu_0 = 0.423$ m^2/s, the major radius $R_0 = 1.69$ m, the mean minor radius $a = 0.65$, the mean ion density $n_i = 5.0 \times 10^{19}$, the magnetic field strength at the magnetic axis ($B_0 \approx 1.84$ T) and the poloidal grid resolution (48 radial by 64 axial meshpoints). Fourth-order polynomial interpolation was used within individual finite elements in the poloidal plane.

References

[1] R.J. La Haye, Neoclassical tearing modes and their control, Phys. Plasmas 13 (2006) 055501.
[2] Z. Chang, J.D. Callen, E.D. Fredrickson, R.V. Budny, C.C. Hegna, K.M. McGuire, M.C. Zarnstorff, TFTR group, Observation of nonlinear neoclassical pressure-gradient-driven tearing modes in TFTR, Phys. Rev. Lett. 74 (1995) 4663.
[3] Z. Chang, E. Fredrickson, J. Callen, K. McGuire, M. Bell, R. Budny, C. Bush, D. Darrow, A. Janos, L. Johnson, H. Park, S. Scott, J. Strachan, E. Synakowski, G. Taylor, R. Wieland, M. Zarnstorff, S. Zweben, Transport effects of low (m,n) MHD modes on TFTR supershots, Nucl. Fusion 34 (1994) 1309.
[4] S. Günter, A. Gude, M. Maraschek, Q. Yu, The ASDEX Upgrade Team, Influence of neoclassical tearing modes on energy confinement, Plasma Phys. Control. Fusion 41 (1999) 767.
[5] E.J. Strait, L. Lao, A.G. Kellman, T.H. Osborne, R. Snider, R.D. Stambaugh, T.S. Taylor, MHD instabilities near the β limit in the doublet III-D tokamak, Phys. Rev. Lett. 62 (1989) 1282–1285.
[6] O. Sauter, R.J. La Haye, Z. Chang, D.A. Gates, Y. Kamada, H. Zohm, A. Bondeson, D. Boucher, J.D. Callen, M.S. Chu, T.A. Gianakon, O. Gruber, R.W. Harvey, C.C. Hegna, L.L. Lao, D.A. Monticello, F. Perkins, A. Pletzer, A.H. Reiman, M. Rosenbluth, E.J. Strait, T.S. Taylor, A.D. Turnbull, F. Waelbroeck, J.C. Wesley, H.R. Wilson, R. Yoshino, Beta limits in long-pulse tokamak discharges, Phys. Plasmas 4 (1997) 1654–1664.
[7] R. La Haye, L. Lao, E. Strait, T. Taylor, High beta tokamak operation in DIII-D limited at low density/collisionality by resistive tearing modes, Nucl. Fusion 37 (1997) 397.
[8] ITER Physics Basis Editors, ITER Physics Expert Group Chairs and Co-Chairs, ITER Joint Central Team and Physics Integration Unit, Chapter 1: overview and summary, Nucl. Fusion 39 (1999) 2137.
[9] H. Zohm, G. Gantenbein, G. Giruzzi, S. Günter, F. Leuterer, M. Maraschek, J.P. Meskat, A.G. Peeters, W. Suttrop, D. Wagner, Experiments on neoclassical tearing mode stabilization by ECCD in ASDEX Upgrade, Nucl. Fusion 39 (1999) 577–580.
[10] G. Gantenbein, H. Zohm, G. Giruzzi, S. Günter, F. Leuterer, M. Maraschek, J.P. Meskat, Q. Yu, A. Team, Complete suppression of neoclassical tearing modes with current drive at the electron-cyclotron-resonance frequency in ASDEX Upgrade tokamak, Phys. Rev. Lett. 85 (2000) 1242–1245.
[11] A. Isayama, Y. Kamada, S. Ide, K. Hamamatsu, T. Oikawa, T. Suzuki, Y. Neyatani, T. Ozeki, Y. Ikeda, K. Kajiwara, The JT-60 Team, Complete stabilization of a tearing mode in steady state high-β_p H-mode discharges by the first harmonic electron cyclotron heating/current drive on JT-60U, Plasma Phys. Control. Fusion 42 (2000) L37.
[12] R.J. La Haye, S. Günter, D.A. Humphreys, J. Lohr, T.C. Luce, M. Maraschek, C.C. Petty, R. Prater, J.T. Scoville, E.J. Strait, Control of neoclassical tearing modes in DIII–D, Phys. Plasmas 9 (2002) 2051.
[13] D.A. Humphreys, J.R. Ferron, R.J. La Haye, T.C. Luce, C.C. Petty, R. Prater, A.S. Welander, Active control for stabilization of neoclassical tearing modes, Phys. Plasmas 13 (2006) 056113.

[14] R. Prater, R. La Haye, T.C. Luce, C.C. Petty, E.J. Strait, J.R. Ferron, D.A. Humphreys, A. Isayama, J. Lohr, K. Nagasaki, P.A. Politzer, M.R. Wade, A.S. Welander, Stabilization and prevention of the 2/1 neoclassical tearing mode for improved performance in DIII-D, Nucl. Fusion 47 (2007) 371–377.
[15] A. Isayama, Y. Kamada, N. Hayashi, T. Suzuki, T. Oikawa, T. Fujita, T. Fukuda, S. Ide, H. Takenaga, K. Ushigusa, T. Ozeki, Y. Ikeda, N. Umeda, H. Yamada, M. Isobe, Y. Narushima, K. Ikeda, S. Sakakibara, K. Yamazaki, K. Nagasaki, The JT-60 Team, Achievement of high fusion triple product, steady-state sustainment and real-time NTM stabilization in high-β ELMy H-mode discharges in JT-60U, Nucl. Fusion 43 (2003) 1272.
[16] M. Maraschek, G. Gantenbein, Q. Yu, H. Zohm, S. Günter, F. Leuterer, A. Manini, Enhancement of the stabilization efficiency of a neoclassical magnetic island by modulated electron cyclotron current drive in the ASDEX Upgrade tokamak, Phys. Rev. Lett. 98 (2007) 1.
[17] C.C. Hegna, J.D. Callen, A closure scheme for modeling RF modifications to the fluid equations, Phys. Plasmas 16 (2009) 112501.
[18] J.J. Ramos, Fluid and drift-kinetic description of a magnetized plasma with low collisionality and slow dynamics orderings. I. Electron theory, Phys. Plasmas 17 (2010) 082502.
[19] J.J. Ramos, Fluid and drift-kinetic description of a magnetized plasma with low collisionality and slow dynamics orderings. II. Ion theory, Phys. Plasmas 18 (2011) 102506.
[20] I.B. Bernstein, L. Friedland, Handbook of Plasma Physics, vol. 1, North-Holland, Amsterdam, 1983, pp. 367–418.
[21] T.G. Jenkins, S.E. Kruger, Fluid equations in the presence of electron cyclotron current drive, Phys. Plasmas 19 (2012) 122508.
[22] C.R. Sovinec, A.H. Glasser, T.A. Gianakon, D.C. Barnes, R.A. Nevel, S.E. Kruger, S.J. Plimpton, A. Tarditi, M.S. Chu, The NIMROD Team, Nonlinear magneto-hydrodynamics with high-order finite elements, J. Comput. Phys. 195 (2004) 355.
[23] A.P. Smirnov, R.W. Harvey, K. Kupfer, A general ray tracing code GENRAY, Bull. Am. Phys. Soc. 39 (1994) 1626.
[24] S.E. Kruger, D.D. Schnack, C.R. Sovinec, Dynamics of the major disruption of a DIII-D plasma, Phys. Plasmas 12 (2005) 056113.
[25] R. Takahashi, D.P. Brennan, C.C. Kim, Kinetic effects of energetic particles on resistive MHD stability, Phys. Rev. Lett. 102 (2009) 135001.
[26] B.J. Burke, S.E. Kruger, C.C. Hegna, P. Zhu, P.B. Snyder, C.R. Sovinec, E.C. Howell, Edge localized linear ideal magnetohydrodynamic instability studies in an extended-magnetohydrodynamic code, Phys. Plasmas 17 (2010) 032103.
[27] V. Izzo, E. Hollmann, A. James, J. Yu, D. Humphreys, L. Lao, P. Parks, P. Sieck, J. Wesley, R. Granetz, G. Olynyk, D. Whyte, Runaway electron confinement modelling for rapid shutdown scenarios in DIII-D, Alcator C-Mod and ITER, Nucl. Fusion 51 (2011) 063032.
[28] J.R. King, C.R. Sovinec, V.V. Mirnov, First-order finite-Larmor-radius effects on magnetic tearing in pinch configurations, Phys. Plasmas 18 (2011) 042303.
[29] E.D. Held, J.D. Callen, C.C. Hegna, C.R. Sovinec, T.A. Gianakon, S.E. Kruger, Nonlocal closures for plasma fluid simulations, Phys. Plasmas 11 (2004) 2419.
[30] P.T. Bonoli, J. Ko, R. Parker, A.E. Schmidt, G. Wallace, J.C. Wright, C.L. Fiore, A.E. Hubbard, J. Irby, E. Marmar, M. Porkolab, D. Terry, S.M. Wolfe, S.J. Wukitch, The Alcator C-Mod Team, J.R. Wilson, S. Scott, E. Valeo, C.K. Phillips, R.W. Harvey, Lower hybrid current drive experiments on Alcator C-Mod: comparison with theory and simulation, Phys. Plasmas 15 (2008) 056117.
[31] A. Schmidt, P.T. Bonoli, O. Meneghini, R.R. Parker, M. Porkolab, S. Shiraiwa, G. Wallace, J.C. Wright, R.W. Harvey, J.R. Wilson, Investigation of lower hybrid physics through power modulation experiments on Alcator C-Mod, Phys. Plasmas 18 (2011) 056122.
[32] C.B. Forest, P.K. Chattopadhyay, R.W. Harvey, A.P. Smirnov, Off-midplane launch of electron Bernstein waves for current drive in overdense plasmas, Phys. Plasmas 7 (2000) 1352–1355.
[33] C. Forest, R. Harvey, A. Smirnov, Power deposition by mode converted electron Bernstein waves in the DIII-D 'heat pinch' experiments, Nucl. Fusion 41 (2001) 619.
[34] D.D. Schnack, Lectures in magnetohydrodynamics: with an appendix on extended MHD, Springer, Berlin, Heidelberg, 2009.
[35] J.J. Ramos, Finite-Larmor-radius kinetic theory of a magnetized plasma in the macroscopic flow reference frame, Phys. Plasmas 15 (2008) 082106.
[36] B.D. Fried, S.D. Conte, The Plasma Dispersion Function, Academic Press, New York, NY, 1961.
[37] D.D. Schnack, D.C. Barnes, D.P. Brennan, C.C. Hegna, E. Held, C.C. Kim, S.E. Kruger, A.Y. Pankin, C.R. Sovinec, Computational modeling of fully ionized magnetized plasmas using the fluid approximation, Phys. Plasmas 13 (2006) 058103.
[38] C. Sovinec, J. King, Analysis of a mixed semi-implicit/implicit algorithm for low-frequency two-fluid plasma modeling, J. Comput. Phys. 229 (2010) 5803.
[39] T.G. Jenkins, S.E. Kruger, C.C. Hegna, D.D. Schnack, C.R. Sovinec, Calculating electron cyclotron current drive stabilization of resistive tearing modes in a nonlinear magnetohydrodynamic model, Phys. Plasmas 17 (2010) 012502.
[40] L. Friedland, I.B. Bernstein, Geometric optics in plasmas characterized by non-Hermitian dielectric tensors, Phys. Rev. A 22 (1980) 1680.
[41] C.B. Barber, D.P. Dobkin, H.T. Huhdanpaa, The Quickhull algorithm for convex hulls, ACM Trans. Math. Softw. 22 (1996) 469, see also http://www.qhull.org.
[42] D. Humphreys, J. Ferron, M. Bakhtiari, J. Blair, Y. In, G. Jackson, H. Jhang, R. Johnson, J. Kim, R.L. Haye, J. Leuer, B. Penaflor, E. Schuster, M. Walker, H. Wang, A. Welander, D. Whyte, Development of ITER-relevant plasma control solutions at DIII-D, Nucl. Fusion 47 (2007) 943.
[43] R.J. Renka, Algorithm 660, QSHEP2D: quadratic Shepard method for bivariate interpolation of scattered data, ACM Trans. Math. Softw. 14 (1988) 149.
[44] R.J. Renka, R. Brown, Algorithm 792: accuracy test of ACM algorithms for interpolation of scattered data in the plane, ACM Trans. Math. Softw. 25 (1999) 78.
[45] D. Shepard, A two-dimensional interpolation function for irregularly-spaced data, in: Proceedings of the 23rd ACM National Conference, ACM '68, ACM, New York, NY, USA, 1968, p. 517.
[46] T.G. Jenkins, ECCD-induced tearing mode stabilization via active control in coupled NIMROD/GENRAY HPC simulations, presented orally at International Sherwood Fusion Theory Conference, Santa Fe, NM, April 2013.
[47] N.J. Fisch, A.H. Boozer, Creating an asymmetric plasma resistivity with waves, Phys. Rev. Lett. 45 (1980) 720.
[48] T. Ohkawa, Steady state operation of tokamaks by RF heating, Rep. GA-A-138476, General Atomics, San Diego, CA, 1980, 45 pp.
[49] E.D. Held, S.E. Kruger, J.-Y. Ji, E.A. Belli, B.C. Lyons, Verification of continuum drift kinetic equation solvers in NIMROD, Phys. Plasmas 22 (2015) 032511.
[50] J.-Y. Ji, E.D. Held, Exact linearized Coulomb collision operator in the moment expansion, Phys. Plasmas 13 (2006) 102103.