Princeton Plasma Physics Laboratory

PPPL- 4803

PPPL- 4803

Energy Conservation Tests of a Coupled Kinetic-kinetic Plasma-neutral Transport Code

D.P. Stotler, C.S. Chang, S.H. Ku, J. Lang and G.. Park

August, 2012

PPPL PRINCETON PLASMA PHYSICS LABORATORY

Prepared for the U.S. Department of Energy under Contract DE-AC02-09CH11466.

Princeton Plasma Physics Laboratory
Report Disclaimers

Full Legal Disclaimer

Trademark Disclaimer

PPPL Report Availability

Princeton Plasma Physics Laboratory:

http://www.pppl.gov/techreports.cfm

Office of Scientific and Technical Information (OSTI):

http://www.osti.gov/bridge

Related Links:

U.S. Department of Energy

Office of Scientific and Technical Information

Fusion Links

Energy conservation tests of a coupled kinetic-kinetic plasma-neutral transport code

D P Stotler, C S Chang, G Park[1] and S H Ku

Princeton Plasma Physics Laboratory, Princeton University Princeton, NJ 08543

E-mail: dstotler@pppl.gov

[1]National Fusion Research Institute, Daejeon, Korea

Abstract. An approach to coupling two kinetic particle codes for the simulation of neutral-plasma interactions in magnetic fusion devices is described. The behavior of the neutral atoms and molecules is modeled with a Monte Carlo code. The plasma species are simulated with a particle-in-cell code that integrates the guiding center equations of motion and computes a self-consistent electric field. The coupling algorithm is designed to conserve mass in the neutral-plasma exchanges to statistical accuracy. Although energy is not fully conserved due to the velocity dependence of the charge exchange cross section and the kinetic character of both species, the impact of this non-conservation on the overall simulation is negligible.

PACS numbers: 52.65.Pp, 52.65.Rr, 52.55.Fa, 52.25.Ya, 52.25.Fi

1. Introduction

The design of the ITER magnetic fusion experiment is based on extrapolations of empirical scaling relations for the plasma confinement time developed from present day devices [1]. Since the fusion gain (ratio of fusion to heating power) achieved by ITER is very sensitive to deviations from those scalings, the magnetic confinement fusion community has made a concerted effort to understand the physics underlying the scalings so as to increase our confidence in their extrapolation to the dimensionless parameters anticipated for ITER.

The principal plasma confinement mode targeted for ITER, the so-called "high confinement mode" or "H-mode", is characterized by a substantial barrier to outward radial plasma transport just inside the last closed magnetic flux surface (the magnetic separatrix in a diverted configuration like ITER). The resulting plasma density and temperature profiles possess extremely steep radial gradients just inside the last closed surface, but are much flatter between there and the magnetic axis. Consequently, this steep gradient region is referred to as the "H-mode pedestal".

The associated steep gradients in plasma energy and current are potential sources of free energy for magnetohydrodynamic (MHD) modes of various scales. The most common instabilities in this region, termed "Edge Localized Modes" or ELMs, act to periodically flatten the pedestal gradients. As long as the plasma remains in H-mode, the pedestal builds up again rapidly and the cycle repeats.

The effort to understand the H-mode has proceeded on multiple fronts, attempting to answer key questions:

(i) What physics controls the transition from "low confinement" or "L-mode" into the H-mode? In particular, what determines the observed threshold heating power for this transition?

(ii) What instabilities drive the wide variety of observed ELMs and how can we control them?

(iii) What transport processes govern the pedestal build-up and the subsequent near-equilibrium phase of the H-mode?

The research effort and codes described in this paper are targeted at the third of these questions. Our understanding of plasma transport in magnetic confinement devices is poorest with regard to the anomalous radial transport associated with plasma microturbulence. However, in the H-mode pedestal this microturbulence is much less virulent than in the L-mode, and radial transport fluxes are reduced to levels consistent with "neoclassical" theory. Neoclassical transport sets an upper limit on our ability to confine an assembly of plasma particles orbiting in a toroidal magnetic field and undergoing self-collisions. These effects result in unequal losses of electrons and ions, giving rise to plasma currents and electric fields. Because neoclassical transport theory is a relatively mature topic, an examination of its implications for the H-mode pedestal is an obvious target for making progress towards a greater understanding.

Typical H-mode pedestal parameters, however, violate the assumptions of most analytic neoclassical theories: the pedestal gradient scale lengths can be narrower than the neoclassical ("banana") ion orbit widths, the plasma ion distribution is expected to deviate significantly from a Maxwellian, and the role of the magnetic separatrix, i.e., magnetic topology, is significant. For these, and other reasons that have become apparent over the last several years, accurately simulating neoclassical behavior in the H-mode pedestal requires a kinetic, particle based simulation technique. This is the approach taken here, which builds on the initial study by Chang [2] examining the effects of a self-consistent radial electric field, collisions, and neutral particles on the pedestal characteristics. The need for a kinetic treatment was confirmed by the observation of significant non-Maxwellian ion distributions. Recycled neutral atoms were found to result in a build-up of the pedestal density, provided anomalous radial transport was indeed less than that due to neoclassical effects.

We describe here improvements to the neutral transport routine in the code developed by Chang (XGC0) that yield a more realistic model of plasma-neutral interactions and the ability to incorporate plasma-wall interactions. First, we resolve plasma and neutral quantities poloidally, as well as radially, throughout the entire vacuum vessel volume, allowing the neutral sources to be placed at material surfaces. The spatial variation of those sources is determined directly from the distribution of ions lost to the walls in a physically realistic way. To allow the resulting neutral profile to evolve consistently, the neutral transport is simulated in a time dependent manner and synchronized with the plasma time evolution. Both plasma and neutral species have three-dimensional flow velocity vectors so that the effects of neutral particles on the transport of plasma momentum are treated. The atomic physics interactions between plasma and neutral particles are characterized in detail, including the density dependence of the electron impact ionization of hydrogen and the velocity dependence of the charge exchange cross section. Molecules and impurities will be incorporated in a subsequent version of the code.

The algorithm used for these calculations is described in section 2. The results of a standard test case, comparable to that used in [2], are presented in section 3. The conservation properties of the code in these example runs are examined in section 4, documenting an essential aspect of the verification of this algorithm. Finally, some conclusions and plans for future improvements are described in section 5.

2. Algorithm

The treatment of ion-ion collisions in the XGC0 code [2] is complicated by the nonlinearity of the collision operator. In fact, the development of techniques for accurately and efficiently simulating these collisions remains an active area of research. The collision operator governing plasma-neutral interactions is likewise effectively nonlinear due to the tight coupling between plasma ions and atoms in most regions of the tokamak boundary.

Our approach for dealing with this nonlinearity is closely related to the "test particle Monte Carlo" method attributed to Haviland [3]. In Haviland's original technique, and in that of subsequent practitioners (e.g., [4, 5, 6]), individual test particles collide with a background characterized by a specified distribution function. In the purely nonlinear case, that distribution is typically updated in iterative fashion from the test particle trajectories until a converged (for steady state problems) result is obtained. The accuracy of this method hinges on the adequacy of the chosen representation of the background distribution function, e.g., a drifting Maxwellian.

The next level of sophistication, Direct Simulation Monte Carlo [7], bypasses this limitation by colliding the current test particle with other test particles present in the same computational cell. Because this effectively represents a Monte Carlo discretization of six dimensions of velocity space, versus three for test particle Monte Carlo, the computational resources required to achieve a specified precision in the results are significantly greater. While we do not consider this approach in the present work, the continued and rapid increase in available computational speed will eventually render it practical.

A significant difference between our problem and those addressed in the aforementioned references is that the transport of the plasma (i.e., charged) and neutral species is fundamentally different in character, and the two phenomena have frequently been simulated with separate codes. The most familiar examples to magnetic fusion researchers are coupled fluid plasma and Monte Carlo neutral transport, e.g., B2-EIRENE [8, 9] and UEDGE-DEGAS2 [10]. The coupling between such codes has largely been explicit and iterative or time dependent. The two-dimensional plasma density, flow velocity, and temperature fields produced by the plasma code are input as a background to the neutral transport code; the plasma fluxes to the material boundary provide the source of neutral particles via plasma-material interactions, referred to as "recycling" [11]. Execution of the neutral transport code yields volumetric sources (or sinks) of plasma mass, momentum, and energy that are transferred back to the plasma code for use as the "right hand sides" of its transport equations in its next iteration or time step advance. Adapting this method to the present application is problematic since all of the kinetic details of the plasma-neutral interactions simulated by the neutral routine are lost in the process of compiling the volumetric sources.

We instead follow Chang's [2] approach of extending the test particle Monte Carlo method and employ two complementary plasma-neutral operators, one embedded in the plasma code and one in the neutral transport routine. The neutral collision operator works in a manner very similar to that used in coupling to fluid plasma codes: the kinetic neutral species collide with a plasma background characterized by a particular distribution function, the parameters of which are specified by the plasma code. Instead of returning volumetric sources, the neutral routine provides the parameters needed to specify the background neutral distribution in the plasma code's collision operator.

The accuracy and realism of the overall simulation hinges on the two operators being consistent with each other and the representation of the distributions being appropriate

to the problem at hand. In the present work, the plasma transport calculation is performed by the XGC0 code originally described by Chang [2] and subsequently extended by others [12]; the neutral transport routine is based on the DEGAS2 Monte Carlo neutral transport code [13]. Henceforth, we will refer to the implementations of these two complementary plasma-neutral collision operators as the "XGC0" (kinetic plasma particles colliding on a neutral background) and "DEGAS2" (kinetic neutral particles colliding on a plasma background) collision routines.

The XGC0 code reduces the ion phase space from six dimensions to five by simulating only the motion of the guiding centers of the ion orbits, averaging over the much faster gyromotion. That is, the ion velocity is written as

$$\boldsymbol{v} = \boldsymbol{v}_{\rm gc} + v_\perp(-\sin\theta\hat{\boldsymbol{e}}_1 - \cos\theta\hat{\boldsymbol{e}}_2), \tag{1}$$

where $\boldsymbol{v}_{\rm gc}$ is the guiding center velocity. The second term represents the ion gyromotion with speed $v_\perp$, gyrophase θ and unit vectors perpendicular to the local magnetic field, $\hat{\boldsymbol{e}}_1 \times \hat{\boldsymbol{e}}_2 = \hat{\boldsymbol{b}}$, where $\hat{\boldsymbol{b}} = \boldsymbol{B}/B$ is the magnetic field direction and $B = |\boldsymbol{B}|$ is its magnitude.

The Hamiltonian equations of motion [2, 14] for the ions are integrated by XGC0 for a discrete set of marker particles; the ion guiding center distribution function is written as

$$f(\boldsymbol{x}, \boldsymbol{v}) = \sum_{k=1,N} w(k)\delta[\boldsymbol{x} = \boldsymbol{x}(k)]\delta[\mu = \mu(k)]\delta[\rho = \rho(k)], \tag{2}$$

where the sum is over the N marker particles, each of which is assigned a statistical weight $w(k)$. The marker particles are localized in phase space first by the normalized velocity of the ions parallel to the magnetic field, $\rho = mv_\parallel/qB$, with m and q being the ion mass and charge, respectively. The second coordinate is the magnetic moment, defined as $\mu = mv_\perp^2/(2B)$ with $v_\perp$ defined by (1). Each marker particle effectively represents a portion of phase space occupied by a large number of actual particles, as indicated by the statistical weight.

Conservation of the total weight of the XGC0 marker particles, $w_{\rm tot} \equiv \sum_k w(k)$, can be checked by computing, at time t_i, the left hand side of:

$$w_{\rm tot}(t_i) + \sum_{j=1}^{i} L(t_j) - \sum_{j=0}^{i} I(t_j) = w_{\rm tot}(t_0) \tag{3}$$

and comparing with the total weight at the beginning of the run, weight, $w_{\rm tot}(t_0)$. The quantity L represents the ion losses to the material boundary (the "limiter") at each time step, and I represents the sources (or sinks) due to neutral-plasma interactions, e.g., due to ionization of atoms. Expressions analogous to (3) can be written for the energy and for each of the components of the momentum vector.

A strict test of conservation would evaluate I using the actual kinetic collisions in the XGC0 collision routine, I_p. But, the DEGAS2 collision routine also yields its own value for I, which we denote by I_n. Since the neutral distribution produced by DEGAS2 reflects the effects of the source (or sink) I and since this neutral distribution is going to be used as the background for subsequent execution(s) of the XGC0 collision routine,

any inequality between I_p and I_n represents a violation of mass conservation. The magnitude of $|I_p - I_n|/[(|I_p| + |I_n|)/2]$ quantifies any deficiencies in our plasma-neutral interaction model. The amount of error introduced into (3) by a difference between I_p and I_n provides a quantitative measure of its impact on the overall plasma evolution.

Because the two collision routines run independently, conservation errors may not manifest themselves as obvious numerical problems. The physics results produced by such a simulation would be in error, but we would in general be unable to detect that error since we do not have in hand a comparably detailed, but independent, calculation. We can instead only insist that the algorithm be consistent with the expected conservation laws by explicit evaluation of (3) and monitoring the relative values of I_p and I_n.

2.1. Ion Distribution Function

The present implementation of the coupling algorithm utilizes a drifting Maxwellian distribution to characterize the background species in the collision routines. The expressions used to evaluate the parameters of the ion distribution have been constructed so as to ensure that the treatment of charge exchange collisions is equivalent in the two routines. The DEGAS2 collision routine samples candidate charge exchange ions from the ion distribution,

$$f_{\rm i}(\boldsymbol{v}) = \frac{n_{\rm i}}{(2\pi T_{\rm i}/m)^{3/2}} \exp\left[\frac{-m}{2T_{\rm i}}(\boldsymbol{v} - \boldsymbol{v}_{\rm f,i})^2\right], \tag{4}$$

where $n_{\rm i}$, $T_{\rm i}$, and $\boldsymbol{v}_{\rm f,i}$ are the ion density, temperature, and flow velocity.

The ion density is computed in the obvious manner,

$$n_i(i_{\rm tri}) = \frac{1}{V(i_{\rm tri})} \sum_{k \in i_{\rm tri}} w(k), \tag{5}$$

where $i_{\rm tri}$ is the index of a (triangular) mesh cell and V is its volume. The sum is over all marker particles k in cell $i_{\rm tri}$.

We then require that the ion flow velocity used in (4) equal the average of the ion velocity (1) over the marker distribution (2):

$$\boldsymbol{v}_{f,i}(i_{\rm tri}) = \frac{1}{n_i(i_{\rm tri})V(i_{\rm tri})} \sum_{k \in i_{\rm tri}} w(k)\boldsymbol{v}_{gc}(k). \tag{6}$$

The ion temperature can be expressed in terms of the average ion energy, $\langle E \rangle$,

$$T_{\rm i} = \frac{2}{3}\left(\langle E \rangle - \frac{1}{2} m \boldsymbol{v}_{\rm f,i}^2\right). \tag{7}$$

We replace $\langle E \rangle$ with the average energy computed using (2) so that

$$T_i(i_{\rm tri}) = \frac{m}{3}\left[\frac{1}{n_i(i_{\rm tri})V(i_{\rm tri})} \sum_{k \in i_{\rm tri}} w(k)v^2(k) - v_{f,i}^2(i_{\rm tri})\right]. \tag{8}$$

To minimize statistical fluctuations associated with the gyrophase in $\boldsymbol{v}$ (1), we write

$$v^2(k) = 2\mu(k)B + v_{gc}^2(k) + 2v_\perp(k)\left[\boldsymbol{v}_{gc}(k) \cdot \hat{\boldsymbol{e}}_1 \sin\theta(k) + \boldsymbol{v}_{gc}(k) \cdot \hat{\boldsymbol{e}}_2 \cos\theta(k)\right]. \tag{9}$$

in the expression for T_i. The gyrophase terms in this expression will average out upon integration over the distribution function.

2.2. Vacuum Vessel Filling Mesh

The computational mesh used by both collision routines in compiling and exchanging information about the neutral and plasma distributions begins with a specification of the tokamak vacuum vessel as a closed polygon. Also required is a characterization of the magnetic equilibrium for the plasma discharge being simulated. The most common representation of such equilibria is as a set of poloidal magnetic flux values ψ on a regular, rectangular grid spanning the poloidal plane of the tokamak. Such equilibria are usually obtained as output files from the EFIT code [15].

To ensure adequate spatial resolution in the mesh cells adjoining the boundary and to provide a convenient basis for compiling particle and heat fluxes to all segments of the boundary, we first divide it so that no segment is longer than a specified minimum length.

Most of the plasma heat and particle fluxes will be deposited, however, in a narrow (on the order of centimeters wide) band adjacent to the points at which the magnetic separatrix intersects the boundary; these locations are the "strike points" [11]. To adequately resolve these fluxes and to allow for the fact that the strike point locations vary with the equilibrium, we all the user to specify a grid of poloidal magnetic flux values, normalized to the magnetic flux of the separatrix surface $\psi_n \equiv \psi/\psi_\mathrm{sep}$, that can be mapped onto the boundary, further refining it.

This normalized poloidal magnetic flux grid is characterized by the width of the two grid segments closest to and farthest from the separatrix, as well as the total width of the grid; this approach is analogous to that used by [16]. These three parameters allow one to derive the coefficients of a quadratic expression for the individual poloidal magnetic flux values, which can then be mapped onto the boundary by interpolation from the equilibrium. Independent sets of three parameters are utilized for the scrape-off layer ($\psi_n > 1$) and private flux region ($\psi_n < 1$) sides of the separatrix. Any boundary points added in the initial subdivision step that fall within these regions are removed to simplify interpretation of the resulting heat flux profiles. The discretization of the boundary in the divertor region for example runs described in section 3 is shown in figure 1.

This boundary is also used in loading and pushing the XGC0 marker particles to provide consistency throughout the code. A determination of whether a particle is inside or outside the boundary is made using the Jordan Curve Theorem [17]. If a just-advanced particle is found to be outside the boundary, the particular segment crossed is identified by interpolating between the particle's previous and current locations. Its weight and energy are added to the boundary and heat flux arrays, and the particle is removed from the problem.

The extreme disparity between plasma transport parallel and perpendicular to

the magnetic field results in parameters [e.g., (8)] having gradients much larger in the direction perpendicular to magnetic flux surfaces than along them. Adequately resolving these gradients with an arbitrary spatial discretization would require a prohibitively large number of cells. For this reason, fluid edge plasma transport codes [9, 18] are based on quasi-orthogonal meshes in which one coordinate is aligned with magnetic flux surfaces. Such meshes suit our purposes as well, once the modifications described below have been incorporated. We need only ensure that the spacing of the aligned surfaces and the second, quasi-orthogonal, coordinate are fine enough resolve spatial variations in the plasma and neutral parameters. Here, we employ the Carre code [16], originally developed to produce meshes for B2 and used previously in DEGAS2 [19] applications.

Carre's algorithm for establishing and optimizing the quasi-orthogonal surfaces in the vicinity of the divertor targets, i.e.., the "ends" of the magnetic flux surface contours, results in a boundary that differs slightly from the one that was provided on input. To ensure that the Carre mesh resides entirely within the vacuum vessel boundary used elsewhere within XGC0 (e.g., to identify lost particles), the boundary input to Carre is shifted inward by 5 mm. This shift is implemented as a "polygon offset" using the Clipper library [20, 21].

The other three boundaries of the Carre mesh coincide with magnetic flux surfaces. The innermost boundary is a magnetic flux contour located somewhere between the separatrix and the magnetic axis; the particular contour used depends on the physics being studied. The outermost boundary is the "last" (i.e., encompassing the greatest poloidal magnetic flux) magnetic flux surface that provides a continuous connection between the inner and outer divertor targets. The volume between this flux surface and the separatrix provides the most direct connection to the core plasma and is referred to as the "scrape-off layer" (SOL) [11]. The volume between the separatrix legs near the strike points is magnetically isolated from the rest of the plasma; this is the "private flux region" (PFR). The outermost flux surface in the PFR is determined via criteria analogous to those for the SOL boundary.

The resulting mesh fills most of the vacuum vessel volume. We tile the space between it and vessel boundary, as well as the region interior to the mesh's inner boundary, with triangles via the Triangle utility [22]. Note that these triangles also fill the narrow gap remaining between the "offset" (i.e., the bottom of the Carre mesh) and actual boundaries. Each of the Carre mesh quadrilaterals is split into two triangles so that the mesh is comprised entirely of triangles. Finally, the triangles are renumbered by a Hilbert space filling curve to make the particle localization procedure more efficient [14]. The mesh used for the calculations described in the remainder of this paper is depicted in figure 2.

This method for discretizing the vacuum vessel volume allows all "boundary conditions" to be established at material surfaces in a manner consistent with our understanding of the physics of plasma- and neutral-material interactions. Relative to the original XGC0 neutral transport model, an increased number of particles is required obtain adequate statistics because plasma quantities are resolved in two dimensions

instead of being averaged over magnetic flux surfaces. The capability for averaging the plasma moments over closed flux surfaces could be added to the XGC0-DEGAS2 code should it be deemed necessary.

The calculation of integrals over the ion marker distribution in other XGC0 subroutines (e.g., for the ion-ion collision operator) are performed via linear interpolation onto flux surfaces or spatial location in the manner typical of particle-in-cell codes. By instead evaluating (8) as the sum over all simulation particles j in a mesh triangle $i_{\rm tri}$, the result becomes equivalent to the volume integrals used in compiling the output of the DEGAS2 routine. Consequently, the ionization rate computed by the XGC0 collision routine would be *exactly* equal to that produced by the DEGAS2 routine if one were to hold the ions fixed in between the two subroutine calls. The number of actual ionizations processed during a given time interval, i.e., the effective ionization rate, will converge to this rate in the limit of a large number of marker particles or a long time interval. This implies that the two routines should conserve mass, at least in a statistical sense. Even though the marker particles are advanced one time step between the calls to the DEGAS2 and XGC0 collision routines, the ionization rates computed by the two routines are indeed equal to within the statistical error, as will be demonstrated in section 4.

2.3. Plasma Collision Algorithm

The null collision method [23] is used to process plasma-neutral collisions in the XGC0 routine. The principal idea behind this technique is to sample collisions along a particle's trajectory using a constant reaction rate $\nu_{\rm max}$, but do nothing to the particle for $1-\nu(\boldsymbol{v})/\nu_{\rm max}$ of these (null) collisions. The algorithm is exact in the statistical sense and is made particularly simple if $\nu_{\rm max}$ can be taken as the maximum reaction rate over the problem volume. The total reaction rate for an ion marker particle with velocity $\boldsymbol{v}$ is the sum over all pertinent neutral collision processes,

$$\nu(\boldsymbol{v}) = \sum_j \nu_j(\boldsymbol{v}) \equiv n_n \sum_j \langle \sigma v_{\rm rel} \rangle_j(\boldsymbol{v}), \tag{10}$$

where n_n is the local neutral density and $v_{\rm rel} = |\boldsymbol{v} - \boldsymbol{v}_{\rm n}|$) is the relative velocity of the ion and a colliding neutral. The angle brackets indicate an average of the cross section for process j, σ_j, over the neutral distribution function.

Neutral-plasma collisions in XGC0 are sufficiently infrequent that the neutral collision routine is invoked only every $N_{\rm ion}$ time steps, $N_{\rm ion} > 1$. This approach represents an approximation to the original null collision algorithm in that the collisions are processed with the particles at the final $\boldsymbol{x}(t = t_0 + N_{\rm ion}\Delta t)$ and $\boldsymbol{v}(t = t_0 + N_{\rm ion}\Delta t)$ of an uncollided trajectory rather than at a sampled $t_0 \le t_c \le t_0 + N_{\rm ion}\Delta t$. If $N_{\rm ion}\Delta t$ is comparable to or less than the time required for the particle to cross a mesh triangle, this approximation will be very good. A strict upper limit on $N_{\rm ion}$ is that the distance $(N_{\rm ion}\Delta t)v$ be less than the local gradient scale lengths for the neutral parameters.

The collision frequency places a looser, second upper bound on $N_{\rm ion}$: $(N_{\rm ion}\Delta t)\nu_{\rm max} < 1$. When this requirement is violated, marker particles could potentially undergo multiple real collisions during $N_{\rm ion}\Delta t$, implying significant modification of the local ion distribution function in those regions. Consequently, we require $N_{\rm ion}\Delta t\nu_{\rm max} < 1$. Note that our routine *does* allow for the marker particles to experience multiple collisions during the interval $N_{\rm ion}\Delta t$ to handle transients with large $\nu_{\rm max}$, particles with very short sampled collision times, etc. For the problems described in sections 3 and 4, this looser criterion yields $N_{\rm ion} \leq 20$. To obtain reasonable statistics for the conservation analysis of section 4, we use $N_{\rm ion} = 20$. Sensitivity studies suggest that reducing $N_{\rm ion}$ below this value does not alter the simulation results by more than a few percent.

2.4. Treatment of Ionization and Charge Exchange

Only deuterium ions and atoms are simulated in this initial version of the coupled plasma-neutral code. As noted previously, the relevant interactions are ionization and charge exchange. The ions are used as a quasi-neutral proxy for the electrons in simulating ionization. That is, ionization is processed using the ion locations and a fixed electron temperature profile specified on input. The electron temperature and density dependent ionization rate is computed using DEGAS2's atomic physics routines, which in turn rely on data obtained from a collisional radiative model [24, 25]. The ionization rate associated with a given ion thus, depends on its location in the problem, but not its velocity.

The charge exchange cross section, in contrast, depends only on the relative velocity of an ion and colliding neutral. The probability for a given ion undergoing a charge exchange collision is based on the corresponding Maxwellian averaged reaction rate $\nu_{\rm cx}$, as in (10). This rate is a function of the local neutral temperature and the ion marker particle's velocity relative to the neutral flow velocity. We randomly sample a gyrophase for the ion marker particle and use (1) to obtain its velocity when evaluating this reaction rate.

The velocity dependence of the cross section is enforced by sampling the colliding neutral via a rejection technique, as depicted in figure 3. A neutral velocity $\boldsymbol{v}_n$ is sampled from the local (Maxwellian) distribution, and the charge exchange cross section $\sigma_{\rm cx}(v_{\rm rel})$, is computed. A uniform deviate $\xi_{\rm cx}$ is obtained and the sampled neutral accepted if $\xi_{\rm cx} < \sigma_{\rm cx}v_{\rm rel}/\max(\sigma_{\rm cx}v_{\rm rel})$, where $\max(\sigma_{\rm cx}v_{\rm rel})$ is the maximum value of $\sigma_{\rm cx}v_{\rm rel}$ for all values of $v_{\rm rel}$.

The charge exchange cross section and rate data are accessed via DEGAS2's atomic physics routines, which in turn interpolate from the original data computed by Krstic and Schultz [26] for quantum mechanically "indistinguishable" particles. A common misconception is that charge exchange and elastic scattering must be simulated separately for a system such as $D^+ + D$. However, as is explained in [26], elastic scattering and charge exchange represent the small and large angle scattering limits, respectively, of the same process. For interaction energies above roughly 1 eV, the

differential scattering cross section peaks strongly near 0 and 180 degrees in the center-of-mass frame. The latter peak is equivalent to classical charge exchange in which the atom and ion effectively swap velocity vectors. The peak near 0 degrees represents elastic scattering. Because those collisions yield small collision angles, the momentum transferred between the ion and atom is negligible and the process can be ignored in a transport simulation. Since the emphasis of the present XGC0-DEGAS2 calculations is on physics close to the separatrix, i.e., well above 1 eV, we take this approach here and employ the "spin exchange" cross section determined by Krstic and Schultz. For energies < 1 eV, the differential scattering cross section is significant for a range of scattering angles, blurring the distinction between "elastic scattering" and "charge exchange". A more complicated sampling procedure would be required in this case [27].

We set the value of $\nu_{\max}$ in the XGC0 collision routine, figure 3, to be the maximum value of the sum of the ionization and charge exchange rates over the entire volume for simplicity. In the problems run to date, the time spent in the plasma-neutral collision routines is too small compared with the rest of the code (see section 2.5) to justify the implementation of a more efficient algorithm, e.g., as in [28].

2.5. Consistent Source and Time Dependence

The principal source of neutral atoms and molecules in most fusion experiments is that due to the recycling of plasma ions (and electrons) at the surrounding material surfaces [11]. Consequently, we base the neutral source on the distribution of ion fluxes to those material surfaces. As noted in section 2, these fluxes are accumulated on a user specified discretization of the simulation boundary. These data are passed on to the DEGAS2 subroutine along with the plasma moments, such as (8).

These fluxes vary in time along with the plasma and neutral profiles, especially during the initial transient phase of the simulations. To allow all of these quantities to evolve consistently, the periodic updates of the neutral profiles are performed in a time dependent manner [29] over the time interval between calls to the DEGAS2 routine. The resulting moments of the neutral distribution function are averaged over that time interval. A time dependent Monte Carlo calculation involves sampling from the recycling source uniformly in time during the interval, as well as from the neutral population that was still in the volume at the end of the previous interval.

This carryover of neutral particles from one time interval to the next also complicates the conservation analysis. To simplify mass conservation on subsequent calls to the XGC0 collision routine, the neutral density profile is multiplied by I_n/I_p where I_n is volume integrated rate of ionization computed by the DEGAS2 routine in its most recent call, and I_p is the volume integrated ionization rate computed with the current ion density. In this way, each execution of the XGC0 collision routine will produce ions at the rate prescribed by the DEGAS2 routine.

To obtain an acceptably accurate (i.e., smooth) poloidal profile for the particle fluxes, we compile them over longer time interval than that between neutral profile

updates (e.g., 1000 vs. 100 steps). At each call to the DEGAS2 subroutine, this profile is rescaled to match the total number of ions lost to the material boundary since the previous call to the DEGAS2 routine.

DEGAS2 can simulate the interactions of plasma ions and neutrals with the surrounding material surfaces with realism limited only by our understanding of those interactions. Utilizing the most detailed available models requires a knowledge of the velocity distribution of the ions as they strike the surface. However, the complicating effects of the sheath and surface structure on the incident angle [30] motivate the use of a simpler model in which a fixed average angle of incidence is assumed and the incident particles are characterized only by their energy distribution. In a future upgrade, machinery for handling the problem boundary (section 2) will be extended to compile this distribution and pass it to the DEGAS2 routine.

The simulations described here instead use an even simpler model in which each wall interaction (of atoms as well as recycling ions) results in a 3 eV atom having a cosine angular distribution relative to the wall normal. The physical basis for this model is that a roughened, saturated surface is expected to recycle a significant fraction of the incident flux as molecules at the wall temperature. The mean free path of such molecules for the plasma conditions of interest is short, on the order of millimeters. The resulting dissociations would yield atoms with Franck-Condon energies of about 3 eV.

The DEGAS2 routine also allows the user to specify the fraction of the incident fluxes that are "recycled" back into the plasma. This recycled fraction may vary along the problem boundary, but is fixed in time. For the simulations reported here, this recycled fraction is 90%. Again, this applies to both recycling ions as well as to neutral atoms striking the boundary along their Monte Carlo trajectory.

3. Sample Results

These example simulations are based on DIII-D discharge 96333 [31]. The EFIT equilibrium for this shot at 3300 milliseconds has served as a standard reference discharge for XGC simulations[32, 33]. As in those papers, we assume initial H-mode like profiles with a pedestal density of 5×10^{19} m^{-3} and a temperature of 1 keV. However, for the present simulations we employ a somewhat lower electron temperature, figure 4(b), consistent with DIII-D H-mode profiles at this density (e.g., [34]). The only other adjustable parameters in these simulations are the 90% recycling coefficient noted in section 2.5 and a collisionless gyroviscosity coefficient [35] of 5×10^{-2} m^2/s used by XGC0 in pushing its ions.

The ion marker particles are tracked for 20 ion transit times (10^4 time steps, 1.56 ms) until all transients have died off. The density pedestal builds up and the gradients steepen, as in [2]. Note that the ion temperature drops from its initial value of 1 keV at the top of the pedestal since we have, for simplicity, not included a heat source from the core to offset the ion heat loss to the boundary and neutral cooling. More detailed XGC0-DEGAS2 simulations of particular experimental discharges would include an appropriate

heat source, kinetic electrons, impurity species, as well as turbulent diffusion with an experimentally calibrated anomalous diffusion coefficient.

The ion fluxes to the divertor floor are plotted in figure 5 as a function of distance along the boundary. These ion losses are solely due to ion motion along and across the open field lines, as well as to ion orbits intersecting the material boundary, with the latter occurring primarily in the low poloidal field region around the X-point [2].

XGC0 scales very efficiently on massively parallel computers, up to peta-flop levels. However, these simulations are performed on a smaller, local Linux cluster utilizing 32 cores on 8 nodes, a practical amount of computational power given that we are not using the XGC0's kinetic electron capability. About 19 hours are required to process the 10,000 time steps in these runs. All of the neutral related computations, including both the XGC0 and DEGAS2 collision routines, occupied about 20% of the time. In comparison, the routines that advance the ion marker positions and velocities typically consume more than half of the time. A total of 7.6 million ions and 0.96 million neutral particles are tracked through the 17,325 mesh cells in the geometry used for the plasma-neutral calculation. Less than 1 GB of memory is needed on each core.

4. Conservation Properties

To assess the conservation properties of the plasma-neutral coupling algorithm, we compare the global integrals of the neutral-plasma exchange rates computed by the two collision routines (I_p and I_n in section 2). The time variation of the ion particle and energy source rates over the last 75% of the run are shown in figure 6(a); for clarity we exclude the initial period dominated by transient relaxation from the input profiles. Charge exchange results in no net production of ions so that the ion source rate represents only the effects of ionization. As expected from the design of the algorithm (section 2), the source rates from the two collision routines match to within the statistical variations.

Note that the XGC0 rates are computed via the equivalent of a "collision estimator" [25], i.e., based entirely on the actual collisions of the XGC0 marker particles. Since only a small fraction of the XGC0 ions experience such collisions in the time interval $N_{\rm ion}\Delta t$, the resulting rates are relatively noisy. Note, however, that these rates are computed only for diagnostic purposes, and the statistical variations have no impact on the outcome of the simulation. The DEGAS2 rates, in contrast, are compiled with a "track length estimator" [25] in which quantities are integrated along the neutral particle trajectories. This approach yields more precise results than the collision estimator in the low collision frequency limit.

The negative values of the ion energy source rates are expected since, on average, the warmer plasma ions are transferring energy to the cooler neutral atoms via the charge exchange process. Ionization serves as an energy source since it adds new ions, albeit cold ones, to the population. The cooling associated with electron impact excitation and ionization of atoms are not accounted for in these rates since the electron temperature

profile is held fixed.

The principal result of interest in figure 6(a) is that the energy source rates from the two collision routines differ significantly, by $\sim 30\%$ at the end of the run. However, an evaluation of the total energy conservation, (3), shows a difference of only 1%, since the energy "source" due to neutral collisions is relatively small compared with transport and the other processes that determine the integrated plasma energy.

To demonstrate that that the discrepancy in the energy source rates is associated with charge exchange and not ionization, we plot in figure 6(b) the charge exchange rate integrated over the entire plasma (left vertical axis) and the energy "source" only due to charge exchange (right axis). The difference between the XGC0 and DEGAS2 curves in figure 6(b) precisely tracks that in figure 6(a).

To go further, we examine the *expected* rates for the charge exchange process and the associated energy exchange for the given neutral background. First, the charge exchange reaction rate evaluated at the local neutral temperature and flow velocity is:

$$\langle \sigma_{\rm cx} v \rangle_{\rm M,n}(\boldsymbol{v}; T_{\rm n}, \boldsymbol{v}_{f,n}) = \int d\boldsymbol{V} f_{\rm M,n}(\boldsymbol{V}; T_n, \boldsymbol{v}_{f,n}) \sigma_{\rm cx}(|\boldsymbol{v}-\boldsymbol{V}|)|\boldsymbol{v}-\boldsymbol{V}| \tag{11}$$

$$= I_{0,0}(|\boldsymbol{v}-\boldsymbol{v}_{f,n}|, T_n) \tag{12}$$

where

$$I_{\ell,n}(u,T) \equiv \frac{\exp(-u^2/u_t^2)}{\sqrt{\pi} u u_t^{n+1}} \int_0^\infty dv_{\rm rel} v_{\rm rel}^{2+n} \sigma_\ell(v_{\rm rel}) \left\{ \exp\left[\frac{-v_{\rm rel}(v_{\rm rel}-2u)}{u_t^2}\right] \right.$$
$$\left. -(-1)^n \exp\left[\frac{-v_{\rm rel}(v_{\rm rel}+2u)}{u_t^2}\right] \right\}, \tag{13}$$

is an integral that arises in taking the moments of atomic physics cross sections over a Maxwellian distribution function [36]. Only the total cross section, corresponding to $\ell = 0$, is of interest for present purposes.

We can then compute in each mesh triangle the velocity space integral of this rate over the local ion marker distribution, (2),

$$\langle \langle \sigma_{\rm cx} v \rangle_{\rm M,n} \rangle_i = \int d\boldsymbol{v} f_i(\boldsymbol{v}) \langle \sigma_{\rm cx} v \rangle_{\rm M,n}(\boldsymbol{v}, T_{\rm n}, \boldsymbol{v}_{f,n}). \tag{14}$$

The sum of this over all mesh triangles yields the charge exhange rate curve labeled "Expected" in figure 6(b).

We can similarly compute the expected energy exchange rate:

$$\langle \langle \Delta E_{\rm cx} \rangle_{\rm M,n} \rangle_i = \frac{m}{2} \int d\boldsymbol{v} f_i(\boldsymbol{v}) \int d\boldsymbol{V} f_{\rm M,n}(\boldsymbol{V}; T_n, \boldsymbol{v}_{f,n})$$
$$\times \sigma_{\rm cx}(|\boldsymbol{v}-\boldsymbol{V}|)|\boldsymbol{v}-\boldsymbol{V}|(V^2 - v^2). \tag{15}$$

That the "Expected" curves track the "XGC0" curves implies that the null collision algorithm is working properly and that the ions are undergoing collisions at the expected rate.

If we instead replace f_i with a Maxwellian distribution at the local T_i and $\boldsymbol{v}_{f,i}$, we get

$$\langle \langle \sigma_{\rm cx} v \rangle_{\rm M,n} \rangle_{\rm M,i} = I_{0,0}[w_f, (T_{\rm i} + T_{\rm n})], \tag{16}$$

where $w_f \equiv |\boldsymbol{v}_{f,i} - \boldsymbol{v}_{f,n}|$, and

$$\langle\langle \Delta E_{\rm cx}\rangle_{\rm M,n}\rangle_{\rm M,i} = \frac{2(\boldsymbol{v}_{f,n} - \boldsymbol{v}_{f,i})}{w_f^2} \cdot (\boldsymbol{v}_{f,i} T_n + \boldsymbol{v}_{f,n} T_i) \times \left(\frac{w_f}{\sqrt{2(T_n + T_i)/m}} I_{0,1} - \frac{I_{0,0}}{2}\right) + (T_n - T_p) I_{0,2}, \tag{17}$$

where all of the $I_{\ell,n}$ are evaluated as $I_{\ell,n}[w_f, (T_{\rm i} + T_{\rm n})]$. The corresponding curves in figure 6(b) are labeled "Maxwellian".

The "Maxwellian" charge exchange and energy source rates differ from both of the corresponding XGC0 and DEGAS2 curves, suggesting that the non-Maxwellian character of both atoms and ions is significant.

Such kinetic effects enter into these rates via the the velocity dependence of the charge exchange cross section. To demonstrate this explicitly, we repeat the simulation with a cross section $\sigma_{\rm cx} \propto 1/v_{\rm rel}$ so that the kinetic reaction rate $\sigma_{\rm cx} v_{\rm rel}$ is independent of the relative particle velocity. In particular, we set $\sigma_{\rm cx} v_{\rm rel} = 3.8 \times 10^{-14}$ m^{-3}/s. The resulting energy source rates for the XGC0 and DEGAS2 routines, shown in figure 7, now match.

We anticipate that the momentum exchange rates in these simulations exhibit discrepancies similar to that described above for the energy exchange. However, they cannot be discerned from the analogous plots due to greater noise levels.

5. Conclusions

We have described a new coupled kinetic-kinetic neutral-plasma transport code that can be used to address critical problems in the edge of tokamak plasmas. An investigation into the conservation properties of the approach used to couple the codes concluded that the energy was not being conserved in neutral-plasma exchanges by as much as 30%. However, these differences represent only about 1% of the global power flow through the system so that we do not expect them to have a significant impact on physics results. Consequently, we are beginning to apply the code and continuing to improve its physics capabilities.

Planned upgrades include compiling the energy distribution of ions lost to the material boundary and using that distribution in determining the kinetic character of the neutral source, as described in section 2.5. This will allow molecules to be introduced into the problem in a realistic manner. We also plan to incorporate XGC0's kinetic electrons [12] into the XGC0 collision routine by using their temperature to determine the ionization rate; the electron cooling associated with the excitation and ionization of atoms can then be accounted for in the simulations, a significant effect. The addition of impurities to the coupled XGC0-DEGAS2 code is also being pursued.

Over a longer time frame, we will investigate ways to improve energy conservation in the XGC0-DEGAS2 coupling. One potential approach would be to employ more sophisticated representations of the distribution functions, e.g., spline fits in velocity

space.

Acknowledgments

This work is supported by U.S. DOE Contract DE-AC02-09CH11466.

References

[1] ITER Physics Basis Authors. Iter physics basis. *Nucl. Fusion*, 39:2137, 1999.

[2] C. S. Chang, Seunghoe Ku, and H. Weitzner. Numerical study of neoclassical plasma pedestal in a tokamak geometry. *Phys. Plasmas*, 11:2649–2667, 2004.

[3] J. K. Haviland and M. L. Lavin. Application of the monte carlo method to heat transfer in a rarified gas. *Phys. Fluids*, 5:1399–1405, 1962.

[4] Thomas W. Tuer and George S. Springer. A test particle monte carlo method. *Computers & Fluids*, 1:399–417, 1973.

[5] S. Chatain, C. Gonella, and P. Roblin. Experimental results and monte carlo simulations for a gadolinium atomic beam produced with a focused electron gun. *J. Phys. D: Appl. Phys.*, 30:360–367, 1997.

[6] Alexander Vikhansky and Markus Kraft. Single-particle method for stochastic simulation of coagulation processes. *Chem. Eng. Sci.*, 60:963–967, 2005.

[7] *Molecular Gas Dynamics*. Clarendon Press, Oxford, 1976.

[8] D. Reiter. Progress in two-dimensional edge modelling. *J. Nucl. Mater.*, 196–198:80–89, 1992.

[9] R. Schneider, X. Bonnin, K. Borrass, D. P. Coster, H. Kastelewicz, D. Reiter, V. A. Rozhansky, and B. J. Braams. Plasma edge physics with b2-eirene. *Contrib. Plasma Phys.*, 46:3–191, 2006.

[10] D. P. Stotler, C. F. F. Karney, M. E. Rensink, and T. D. Rognlien. Coupling of parallelized degas 2 and uedge codes. *Contrib. Plasma Phys.*, 40:221–226, 2000.

[11] P. C. Stangeby. *The Plasma Boundary of Magnetic Fusion Devices*. Institute of Physics Publishing, Philadelphia, 2000.

[12] G. Park, C. S. Chang, I. Joseph, and R. A. Moyer. Plasma transport in stochastic magnetic field caused by vacuum resonant magnetic perturbations at diverted tokamak edge. *Phys. Plasmas*, 17:102503, 2010.

[13] D. P. Stotler and C. F. F. Karney. Neutral transport modeling with degas 2. *Contrib. Plasma Phys.*, 34:392–397, 1994.

[14] M. F. Adams, S.-H. Ku, P. Worley, E. D'Azevedo, J. C. Cummings, and C. S. Chang. Scaling to 150k cores: recent algorithm and performance engineering developments enabling xgc1 to run at scale. *Journal of Physics: Conference Series*, 180:012036, 2009.

[15] L. L. Lao, H. St. John, R. D. Stambaugh, A. G. Kellmann, and W. Pfeiffer. Unknown. *Nucl. Fusion*, 25:1611, 1985.

[16] R. Marchand and M. Dumberry. Carre: a quasi-orthogonal mesh generator for 2d edge plasma modelling. *Comput. Phys. Comm.*, 96:232–246, 1996.

[17] E. Haines. Point in polygon strategies. In Paul Heckbert, editor, *Graphics Gems IV*, pages 24–46. Academic Press, 1994.

[18] T. D. Rognlien, P. N. Brown, R. B. Campbell, T. B. Kaiser, D. A. Knoll, P. R. McHugh, G. D. Porter, M. E. Rensink, and G. R. Smith. 2-d fluid transport simulations of gaseous / radiative divertors. *Contrib. Plasma Phys.*, 34:362–367, 1994.

[19] D. P. Stotler, J. Boedo, B. LeBlanc, R. J. Maqueda, and S. J. Zweben. Progress towards the validation of models of the behavior of neutral helium in gas puff imaging experiments. *J. Nucl. Mater.*, 363–365:688–692, 2007.

[20] X. Chen and S. McMains. Polygon offsetting by computing winding numbers. In *ASME 2005 International Design Engineering Technical Conferences and Computers and Information in*

Engineering Conference (IDETC/CIE2005), number DETC2005-85513, pages 565–575, Long Beach, California, USA, September 24–28, 2005. Design Engineering Division and Computers and Information in Engineering Division.

[21] A. Johnson. Clipper - an open source freeware polygon clipping library, 2012. http://www.angusj.com/delphi/clipper.php.

[22] *Applied Computational Geometry: Towards Geometric Engineering*, volume 1148, page 203. Springer, New York, 1996.

[23] H. R. Skullerud. The stochastic computer simulation of ion motion in a gas subjected to a constant electric field. *J. Phys. D: Appl. Phys.*, 1:1567–1568, 1968.

[24] S. D. Loch, C. P. Ballance, M. S. Pindzola, and D. P. Stotler. The role of excited state ionization data on h and he generalized collisionalradiative coefficients. *Plasma Phys. Control. Fusion*, 51:105006, 2009.

[25] Degas 2 user's manual. http://w3.pppl.gov/degas2/Doc/degas2_all.pdf.

[26] P. S. Krstic and D. R. Schultz. Elastic and related transport cross sections for collisions among isotopomers of $h^+ + h$, $h^+ + h_2$, $h^+ + he$, and $h + h_2$. *At. Plasma-Mater. Data Fus.*, 8:1, 1998. Supplement to the journal Nucl. Fus.

[27] R. J. Kanzleiter, D. P. Stotler, C. F. F. Karney, and D. Steiner. Improved elastic collision modeling in degas 2 for low-temperature plasmas. *Phys. Plasmas*, 8:5064–5069, 2000.

[28] D. Tskhakaya, S. Kuhn, Y. Tomita, K. Matyash, R. Schneider, and F. Taccogna. Self-consistent simulations of the plasma-wall transition layer. *Contrib. Plasma Phys.*, 48:121–125, 2008.

[29] D. Reiter, Chr. May, D. Coster, and R. Schneider. Time dependent neutral gas transport in tokamak edge plasmas. *J. Nucl. Mater.*, 220–222:987–992, 1995.

[30] J. N. Brooks. Modeling of sputtering erosion/redepositionstatus and implications for fusion design. *Fus. Eng. Des.*, 60:515–526, 2002.

[31] L. W. Owen, R. J. Colchin, R. Maingi, M. E. Fenstermacher, T. N. Carlstrom, and R. J. Groebner. Origins and spatial distributions of core fueling in the diii-d tokamak. *J. Nucl. Mater.*, 290–293:464–468, 2001.

[32] X. Q. Xu, E. Belli, K. Bodi, J. Candy, C. S. Chang, R. H. Cohen, P. Colella, A. M. Dimits, M. R. Dorr, Z. Gao, J. A. Hittinger, S. Ko, S. Krasheninnikov, G. R. McKee, W. M. Nevins, T. D. Rognlien, P. B. Snyder, J. Suh, and M. V. Umansky. Dynamics of kinetic geodesic-acoustic modes and the radial electric field in tokamak neoclassical plasmas. *Nucl. Fusion*, 49:065023, 2009.

[33] Weigang Wan, Scott E. Parker, Yang Chen, Gun-Young Park, Choong-Seock Chang, and Daren Stotler. The pinch of cold ions from recycling in the tokamak edge pedestal. *Phys. Plasmas*, 18:056116, 2011.

[34] A. W. Leonard, J. A. Boedo, M. Groth, B. L. Lipschultz, G. D. Porter, D. L. Rudakov, and D. G. Whyte. Particle flux and radial profiles in the sol of diii-d during elming h-mode. *J. Nucl. Mater.*, 363–365:1066–1070, 2007.

[35] S. I. Braginskii. *Reviews of Plasma Physics*, volume 1, chapter Transport processes in a plasma, pages 205–311. Consultants Bureau, New York, 1965.

[36] D. Reiter. Neutral gas transport in fusion devices: Atomic and surface data aspects. In R. K. Janev and H. W. Drawin, editors, *Atomic and Plasma-Material Interaction Processes in Controlled Thermonuclear Fusion*, pages 243–266. Elsevier Science Publishers B.V., New York, 1993.

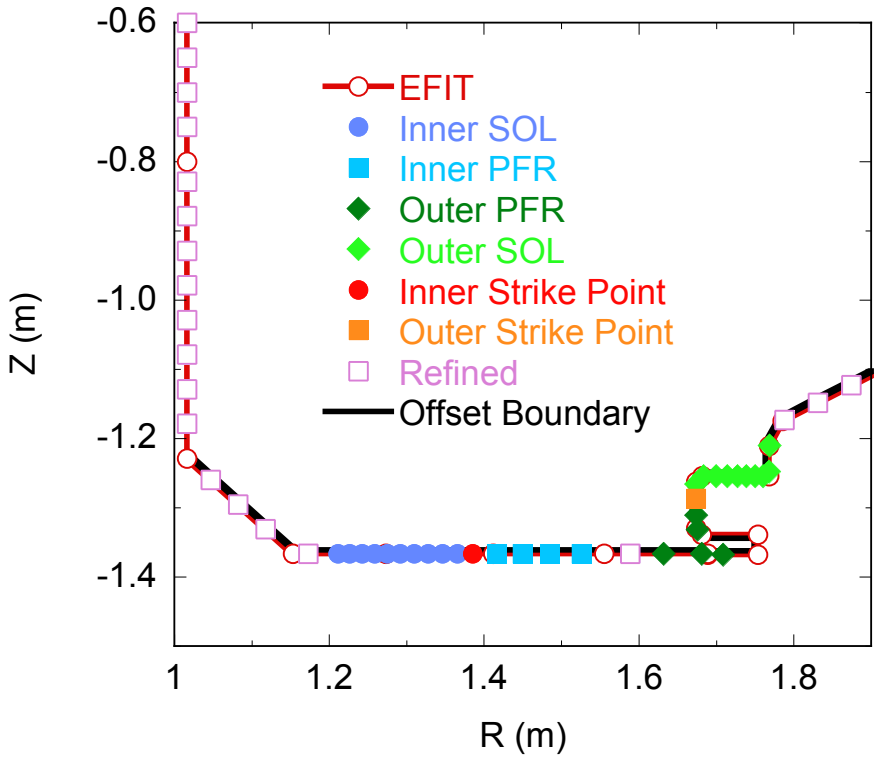

Figure 1. Discretization of the vacuum vessel in the vicinity of the divertor strike points. The points labeled "EFIT" originate with the input magnetic equilibrium. The "Refined" points are added so that no boundary segment is longer than a specified distance, 5 cm here. The magnetic separatrix crosses the boundary at the two "Strike Points". Flux surface based discretizations of the boundary are established on either side of these points. "Inner" and "Outer" refer to smaller and larger values of the major radius R, respectively. The "Offset Boundary" is shifted from the EFIT boundary by 5 mm.

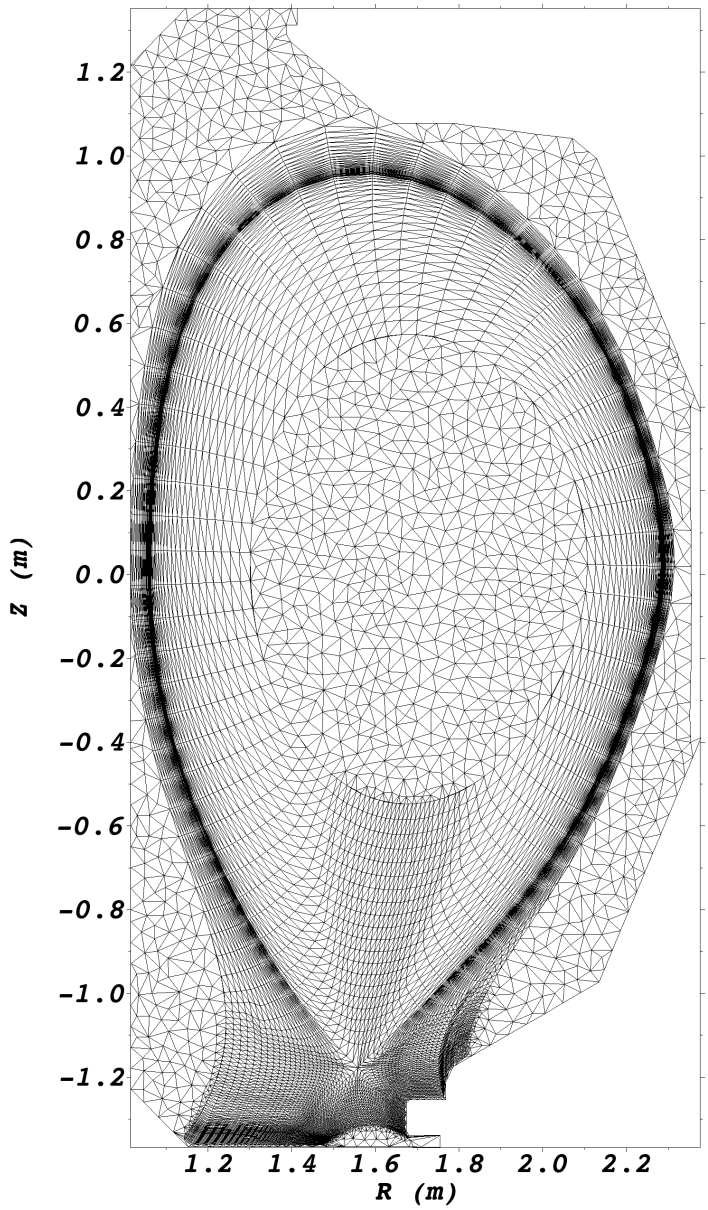

Figure 2. Final triangular mesh used in the plasma-neutral collision routines.

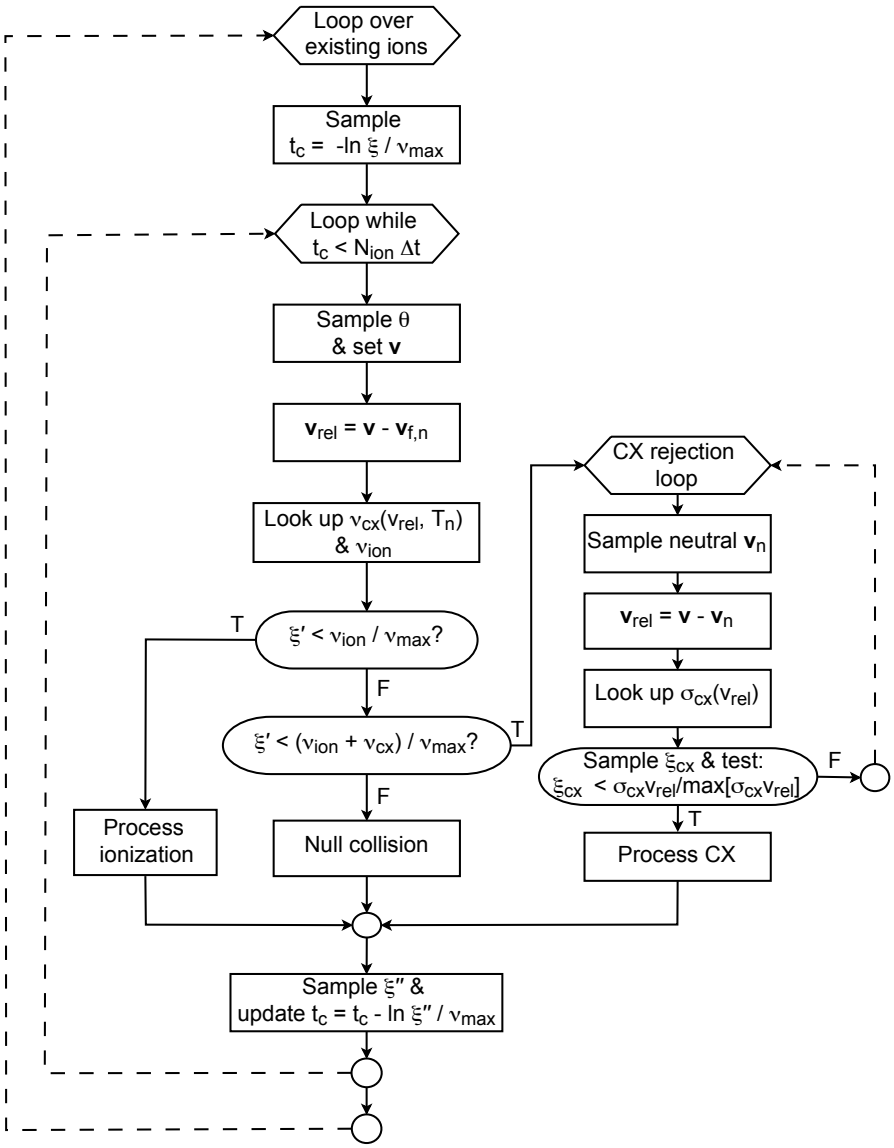

Figure 3. Flowchart for processing ion collisions with a neutral background. The various ξ represent uniform random deviates; all other quantities are defined in the text.

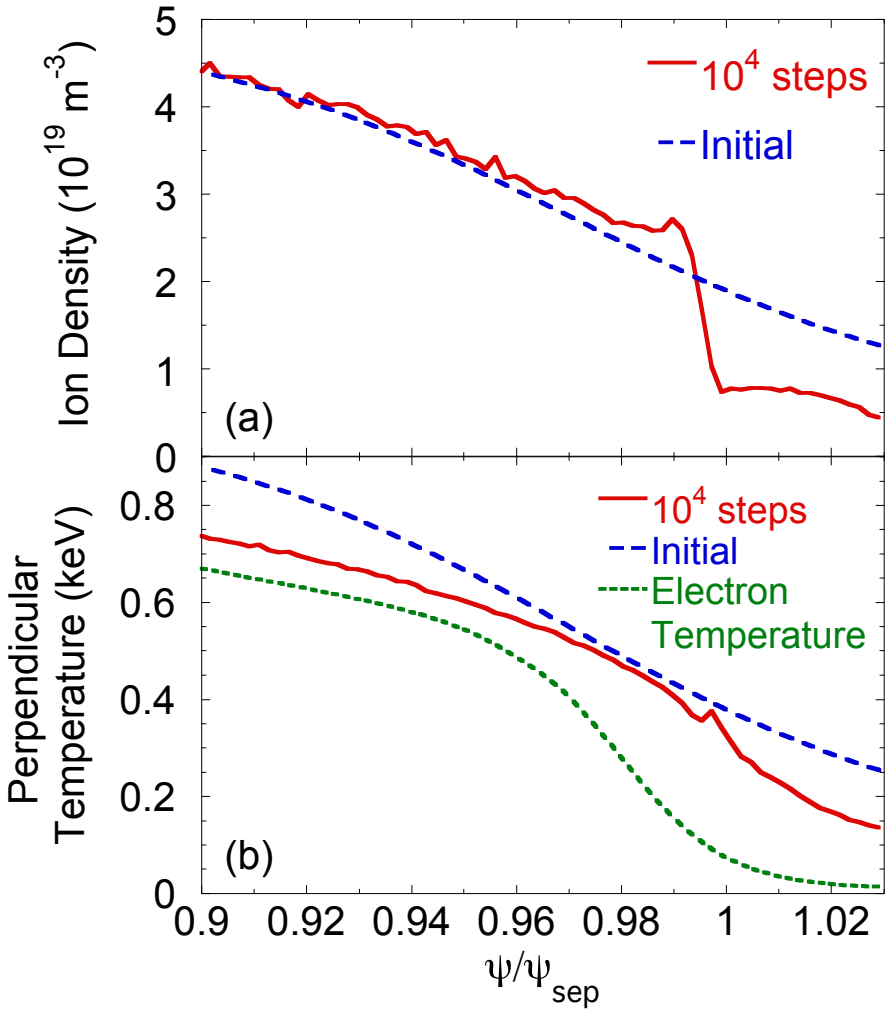

Figure 4. Flux surface averaged ion density (a) and temperature (b) profiles as input and at the end of the simulation. The input electron temperature profile is also shown. The step gradient region in the final ion density profile represents the H-mode pedestal.

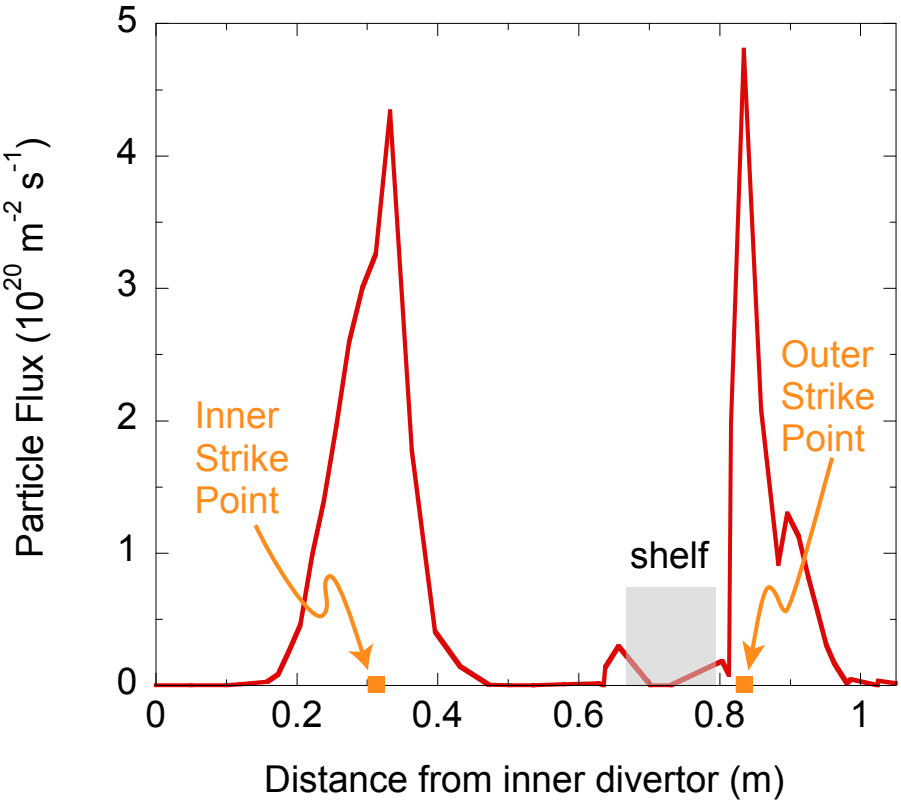

Figure 5. Particle flux to the divertor floor as a function of distance along the boundary from the the inner divertor ($R \simeq 1.2$ m). The shaded region corresponds to the "shelf" region of DIII-D's outer divertor that is shielded from the core plasma.

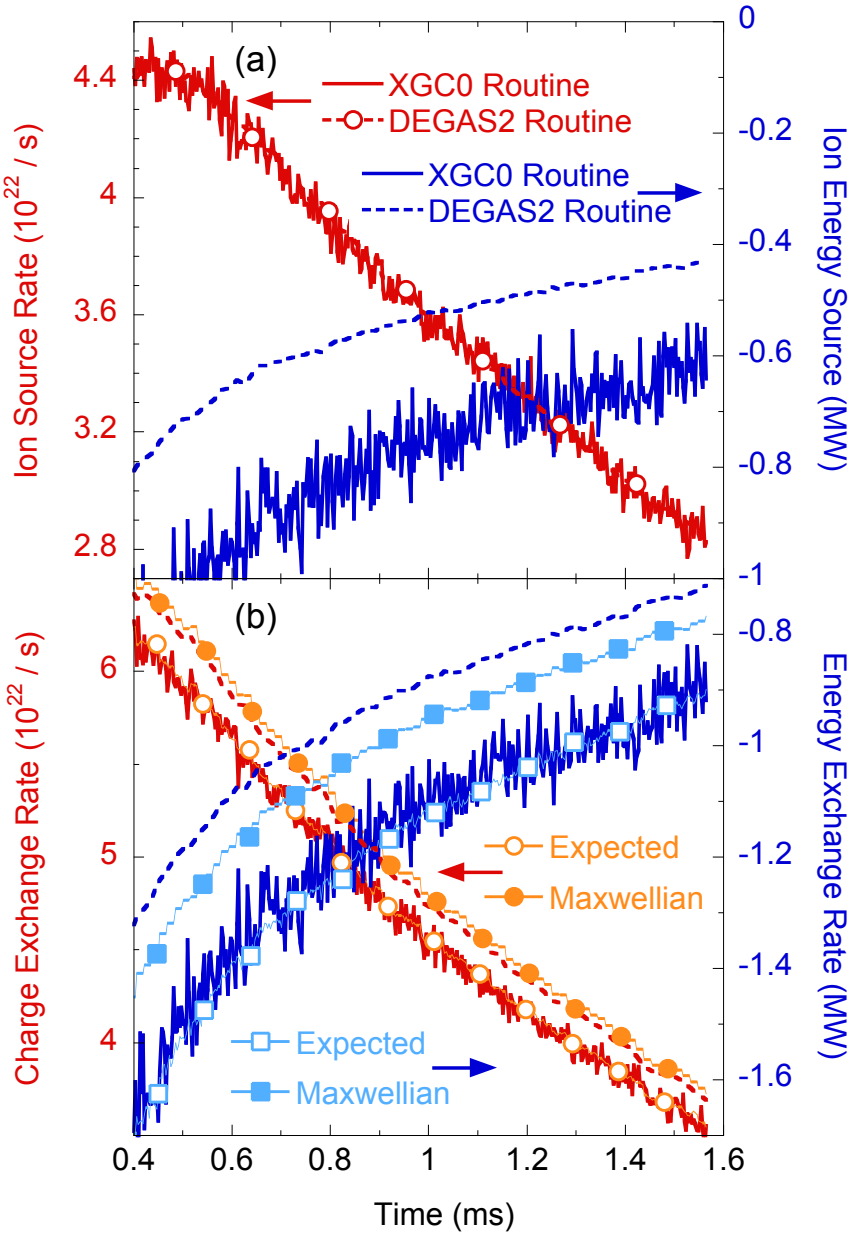

Figure 6. (a) Integrated ion particle (left axis) and energy (right axis) source rates for the two collision routines. The time axis has been truncated to focus on the transport time scale evolution following the initial transients. (b) Corresponding integrated charge exchange (left axis) and energy exchange due to charge exchange rates (right axis) for the two routines. The "Expected" rates are obtained by evaluating the charge exchange rate at the local neutral density and flow velocity, computing the velocity space integral over the ion distribution (2), and summing over all mesh triangles. The "Maxwellian" rates are computed in similar fashion, but utilize the Maxwellian representation of the ion distribution at the local ion temperature and flow velocity.

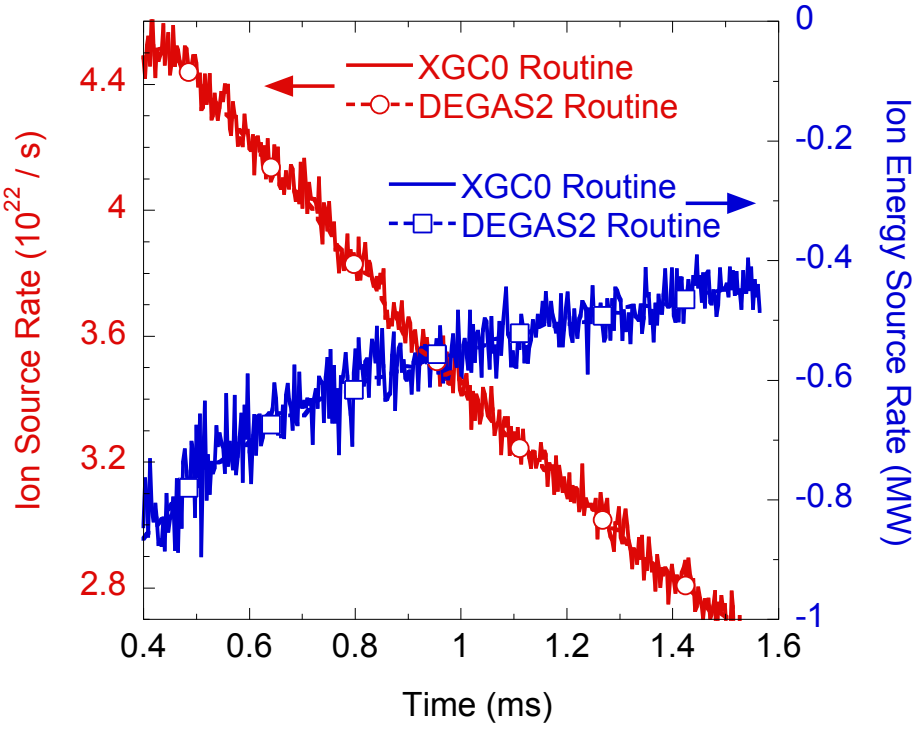

Figure 7. Integrated ion particle (left axis) and energy (right axis) source rates for the two collision routines in a simulation based on a velocity independent charge exchange cross section. The time axis has been truncated to focus on the transport time scale evolution following the initial transients.

The Princeton Plasma Physics Laboratory is operated
by Princeton University under contract
with the U.S. Department of Energy.

Information Services
Princeton Plasma Physics Laboratory
P.O. Box 451
Princeton, NJ 08543

Phone: 609-243-2245
Fax: 609-243-2751
e-mail: pppl_info@pppl.gov
Internet Address: http://www.pppl.gov