This is an electronic reprint of the original article.

This reprint may differ from the original in pagination and typographic detail.

Van Uytven, W.; Dekeyser, W.; Blommaert, M.; Horsten, N.; Marandet, Y.; Baelmans, M.

Advanced spatially hybrid fluid-kinetic modelling of plasma-edge neutrals and application to ITER case using SOLPS-ITER

Published in:

Contributions to Plasma Physics

DOI:

[10.1002/ctpp.202100191](https://doi.org/10.1002/ctpp.202100191)

Published: 01/06/2022

Document Version

Peer reviewed version

Please cite the original version:

Van Uytven, W., Dekeyser, W., Blommaert, M., Horsten, N., Marandet, Y., & Baelmans, M. (2022). Advanced spatially hybrid fluid-kinetic modelling of plasma-edge neutrals and application to ITER case using SOLPS-ITER. *Contributions to Plasma Physics*, 62(5-6), [202100191]. <https://doi.org/10.1002/ctpp.202100191>

Advanced spatially hybrid fluid-kinetic modeling of plasma-edge neutrals and application to ITER case using SOLPS-ITER

W. Van Uytven^{*1}, W. Dekeyser¹, M. Blommaert¹, N. Horsten², Y. Marandet³, and M. Baelmans¹

¹Department of Mechanical Engineering, Ku Leuven, Leuven, Belgium

²Department of Applied Physics, Aalto University, Espoo, Finland

³Aix-Marseille University, CNRS, PIIM, Marseille, France

Abstract

A spatially hybrid fluid-kinetic approach for the efficient simulation of plasma-edge neutrals is extended with a kinetic-fluid condensation process and an automated approach to select a fluid or kinetic treatment at each point along the divertor targets. With the proposed hybrid fluid-kinetic methods, the dominant charge-exchange reaction can be largely removed from the kinetic neutral simulation. The hybrid methods achieve quantitative agreement with the standard kinetic approach, with relative errors on the peak target plasma density and peak target ion flux density of 5-30%, while the CPU time for the kinetic neutral solver is reduced by a factor 3 to 6.

Keywords: plasma edge, hybrid fluid-kinetic neutrals, neutral molecules, SOLPS-ITER, ITER

1 Introduction

The neutral atoms and molecules in plasma-edge simulations are typically modeled with a kinetic Monte Carlo (MC) approach. The computational cost of the MC treatment can, however, become prohibitively large in (partially) detached regimes, mainly due to high charge-exchange (CX) rates between ions and neutral atoms in the so-called high-collisional regions (HCRs). In recent years, significant progress has been made on the development of so-called advanced fluid neutral (AFN) models, which offer an accurate, noise-free, and computationally efficient alternative for the MC treatment of neutral atoms [16, 7]. At the same time, it is well known that this fluid approach for the neutrals is only a good approximation in the HCRs, but not in the entire plasma edge, and not even in the entire divertor. Therefore, different hybrid fluid-kinetic neutral models are being pursued, which aim to combine the benefits of a kinetic and fluid approach in an efficient way [10, 1, 6, 13, 8]. In the standard spatially hybrid models developed in Refs. [1, 8], the sources of neutral atoms are split in kinetic and fluid parts. The division is based on a prior estimate of the validity of the fluid closure in the different plasma regions. In practice this typically means that neutral atoms born at the divertor targets and through volume recombination are treated with a fluid closure, while the kinetic MC treatment is selected for the neutral atoms born at plasma-void interfaces. These are the interfaces between the finite-volume (FV) grid on which the plasma and fluid neutral equations are solved, and the so-called voids, which are regions where only kinetic neutrals are treated, as shown in Figure 1. Neutral molecules are treated kinetically everywhere.

In previous work on the spatially hybrid methods, two causes of a sub-optimal trade-off between speed-up and accuracy were identified. First, kinetic atoms from plasma-void interfaces or dissociating molecules always stayed kinetic until ionization. This meant that kinetic neutrals could still end up undergoing a large amount of collisions in HCRs. Second, the boundary conditions for the neutral atoms were treated with a fluid approach everywhere along the target. However, in realistic geometries, such as that of ITER, the validity of the fluid approach for the atoms varies drastically along the target, being valid in the vicinity of the strikepoint and invalid radially outward.

Figure 1: Illustration of the ITER grid used in this study. The FV-grid, on which the plasma and fluid neutrals are calculated, is shown in green. The blue regions between the FV-grid and the vessel walls are the so-called void regions where only kinetic neutrals are modeled.

This paper aims to resolve the remaining shortcomings of the spatially hybrid approach. First, kinetic atoms resulting from plasma-void interfaces, dissociating molecules, or kinetic boundary regions will be transferred back to the fluid population if they enter an HCR, through so-called condensation reactions. Second, the choice to launch kinetic or fluid neutral atoms at any point along the divertor targets is automated, based on the local collisionality. With these improvements, the spatially hybrid approach is shown to provide excellent agreement with the fully kinetic description, while largely eliminating the dominant CX process from the kinetic MC solver. The work is developed in the new unstructured version of the SOLPS-ITER code suite [3], which consists of an unstructured FV solver for the plasma and fluid neutral equations, coupled to an MC solver for the kinetic neutrals (EIRENE [12]).

The paper is structured as follows: Section 2 explains the different hybrid model variants which will be compared in this study. In Section 3, the performance of the proposed models is analyzed on an ITER case. Finally, Section 4 summarizes the conclusions of the work and presents an outlook for future research.

2 Model description

2.1 Spatially hybrid modeling

The steady-state kinetic Boltzmann equation for the phase-space distribution $f_a(\mathbf{r}, \mathbf{v})$ of the neutral atoms (subscript a) can be written as

$$\mathbf{v} \cdot \nabla f_a(\mathbf{r}, \mathbf{v}) = Q_a^c(\mathbf{r}, \mathbf{v}) + Q_a^v(f_a(\mathbf{r}, \mathbf{v})), \quad (1)$$

where $\mathbf{r}$ indicates the position vector and $\mathbf{v}$ indicates the velocity vector. Q_a represents sinks and sources of kinetic atoms due to surface processes, and sources and collisions in the volume. We distinguish between sources Q_a^c which do not depend on $f_a(\mathbf{r}, \mathbf{v})$, such as atoms from plasma recombination, and sources Q_a^v which do depend on $f_a(\mathbf{r}, \mathbf{v})$, such as CX collisions. In spatially hybrid fluid-kinetic models, the linearity (if neutral-neutral collisions are neglected) of the kinetic equation is exploited by splitting it in a kinetic (subscript k) and fluid (subscript fl) part:

$$\mathbf{v} \cdot \nabla f_{a,k}(\mathbf{r}, \mathbf{v}) = Q_{a,k}^c(\mathbf{r}, \mathbf{v}) + Q_{a,k}^v(f_a(\mathbf{r}, \mathbf{v})) \quad (2)$$

and

$$\mathbf{v} \cdot \nabla f_{a,fl}(\mathbf{r}, \mathbf{v}) = Q_{a,fl}^c(\mathbf{r}, \mathbf{v}) + Q_{a,fl}^v(f_a(\mathbf{r}, \mathbf{v})). \quad (3)$$

The kinetic part of the distribution (Eq. 2) is treated with the standard MC procedure. For the fluid part of the distribution (Eq. 3), a fluid closure is performed, leading to convection-diffusion equations for particles (n), parallel momentum (||m) and energy (H¹):

$$\nabla \cdot \Gamma_a^{n,||m,H} = S_{a,fl}^{n,||m,H}. \quad (4)$$

Here, $\Gamma_a^{n,||m,H}$ represents neutral atom fluxes, and $S_{a,fl}^{n,||m,H}$ represents source terms due to interaction of the fluid neutrals with the plasma by means of ionization, recombination, and CX reactions. Expressions for $\Gamma_a^{n,||m,H}$ and $S_{a,fl}^{n,||m,H}$ in the context of the SOLPS-ITER 2D mean-field plasma edge code, are elaborated in Ref. [16]. The fluid closure needed to go from Eq. 3 to Eq. 4 is elaborated in Ref. [2]. The different physical sources of neutral atoms must be allocated to either $Q_{a,k}$ or $Q_{a,fl}$. Only if the sources for $Q_{a,fl}$ are limited to the regions in the plasma edge where the fluid closure for the neutral atoms is approximately valid, then $f_a \approx f_{a,fl} + f_{a,k}$ will hold. Fluid and kinetic neutral atoms co-exist in the plasma domain and the sources transferred from the neutrals to the plasma are the sum of the contributions from the kinetic and fluid part: $S_{pl}^{n,||m,H} = S_{pl,k}^{n,||m,H} + S_{pl,fl}^{n,||m,H}$. In previous spatially hybrid work, there was limited interaction between the fluid and kinetic components, except for fluid neutrals which are converted to kinetic neutrals when they reach a plasma-void interface, as elaborated in Ref. [1]. Once a kinetic atom was launched, it was followed until ionization and never returned to the fluid population.

2.2 Kinetic-fluid condensation

The validity of the fluid approach for the neutral atoms is predicted by the Knudsen number (Kn), which is defined as $\text{Kn} = \frac{\lambda_{CX}}{L}$, where λ_{CX} is the average CX mean-free path of the local neutral atom distribution, and L is a characteristic macroscopic length scale of the problem. The fluid approach is typically assumed to be strictly valid for $\text{Kn} < 0.01$, $\text{Kn} = 0.1$ is regarded as a transition point, and $\text{Kn} = 1$ predicts the onset of dominant kinetic effects. Similar to Kn, we can define $\text{Kn}^p = \frac{\lambda_{CX}^p}{L}$, where λ_{CX}^p is the CX mean-free path of an individual kinetic particle in the MC simulation. In this work, we add the possibility for kinetic atoms to condense to the fluid neutral population in regions with a high CX-rate. After each CX-collision, it is checked whether $\text{Kn}^p < \text{Kn}^t$, where Kn^t is a user-defined transition threshold. If this is the case, the particle tracing procedure is terminated. EIRENE calculates the source terms due to the interactions of the kinetic atom with the plasma before condensation, $S_{pl,k}^{n,||m,H}$. The mass, parallel momentum, and energy (based on the post-collisional velocity) of the terminated particles are then provided as sources for the fluid neutral atom population ($S_{a,cond}^{n,||m,H}$). The kinetic particle is said to be condensed to the fluid population.

These source terms from condensed neutral atoms are now added to the (steady-state) particle, parallel momentum and energy balance of the fluid neutral atoms (Eq. 4), leading to:

$$\nabla \cdot \Gamma_a^{n,||m,H} = S_{a,fl}^{n,||m,H} + S_{a,cond}^{n,||m,H}. \quad (5)$$

This condensation process is based on earlier work from Ref. [13], itself based on Ref. [9], and was recently found to provide a good trade-off between accuracy and speed-up for atom-only cases without voids [14]. An important difference with the earlier work from Ref. [13] is that a so-called evaporation reaction, where fluid neutrals are converted to kinetic neutrals in the plasma volume, is now omitted. This is done to achieve a higher speed-up, and to avoid buffering issues, where neutrals could often transition back and forth from fluid to kinetic in regimes of intermediate collisionality.

An advantage of the proposed hybrid methods is that the user has freedom in selecting the trade-off between speed-up and accuracy. When performing for example large parameter scans, one might select a higher Kn^t to obtain a larger speed-up, and decrease Kn^t when higher accuracy is required.

¹We use the symbol H instead of E because SOLPS-ITER solves an internal energy equation instead of a total one.

2.3 Treatment of molecules

Although fluid models for neutral molecules are being developed [5], the absence of a dominant scattering collision implies that strong quantitative agreement with the kinetic model cannot be guaranteed. Therefore we choose to always treat molecules kinetically and we instead focus on the efficient treatment of atoms resulting from dissociating molecules. In a recent paper by N. Horsten et. al [8], atoms from dissociation were treated either everywhere kinetically (“Hybrid 1”) or everywhere with a fluid approach (“Hybrid 2”). The former proved most accurate, but speed-up was hereby sacrificed due to MC neutral atoms from the dissociating molecules entering the HCRs. Now that kinetic-fluid condensation is included, we choose to always launch the atoms from dissociation kinetically, capturing first-flight effects. If dissociated atoms end up thermalizing with the plasma, they will condense to the fluid population.

2.4 Model variants for target boundary conditions

All hybrid simulations presented in this paper include kinetic molecular processes, kinetic-fluid condensation, a fluid-kinetic transition at plasma-void interfaces, and a fluid treatment of volume recombination. At the divertor targets, however, we propose three different treatments of the boundary conditions as elaborated below. Table 1 summarizes the entire fluid-kinetic split-up of the proposed hybrid methods, including a comparison to model “Hybrid 1” and “Hybrid 2” from Ref. [8].

Kinetic target boundary conditions Before explaining the different variants for the hybrid models, we elaborate the physical processes which take place at the divertor targets in the reference MC simulation. Ions that strike the target are either recycled as neutral atoms (fast recycling), or thermally released as molecules. The probability depends on the velocity and angle of the incident ion, taking into account that the ion is accelerated in the sheath before hitting the target. In the case of fast recycling, the velocity of the resulting atom also depends on the velocity and angle of the incident ion. The molecules which are thermally released follow a cosine law, with energy determined by the wall temperature (which is assumed to be fixed at 0.1 eV at the targets and 0.05 eV for the other walls in the simulations presented here). For incident atoms, an identical process is followed, but without sheath acceleration. In case of atoms, the process is typically called (fast) reflection instead of recycling. To determine the recycling and reflection probabilities, as well as the energies and angles of the recycled/reflected atoms, EIRENE typically uses the TRIM database.

Model HYB-1: fluid atoms, kinetic molecules Atoms resulting from fast recycling and fast reflection are treated as fluid. The remainder is launched as kinetic molecules (for ease of notation we assume here that there is no absorption at the targets, although this is possible in the code). This boundary condition is identical to that of the hybrid models from Ref. [8]. The flux of reflected and recycled fluid neutral atoms ($\Gamma_{a,fl}$) is given by $\Gamma_{a,fl} = R_{fast,i}\Gamma_{i,inc} + R_{fast,a}\Gamma_{a,inc,fl}$, where $\Gamma_{i,inc}$ and $\Gamma_{a,inc,fl}$ are the incident fluxes of ions and fluid neutral atoms. $R_{fast,i}$ and $R_{fast,a}$ are coefficients of fast recycling and fast reflection, respectively, following from the TRIM tables integrated over the incident velocity space and assuming a Maxwellian distribution for the incident ions and atoms (including sheath acceleration for the ions) [7]. The corresponding fluxes of parallel momentum and energy are calculated self-consistently, as elaborated in Ref. [7]. These fluid neutral boundary conditions have been proven to be accurate in HCRs [16, 7, 15, 8]. A separate surface source of thermally released kinetic molecules ($\Gamma_{mol,k}$) now has to be added, with a strength of $\Gamma_{mol,k} = (1 - R_{fast,i})\Gamma_{i,inc} + (1 - R_{fast,a})\Gamma_{a,inc,fl}$.

Model HYB-2: kinetic relaunch In this model, all incident ions and fluid neutrals are relaunched kinetically, allowing to use the exact TRIM tables on the EIRENE side of the code to determine the reflection probabilities. In EIRENE, we have two sources, $\Gamma_{a,fl \rightarrow a+mol,k}$ (with source strength $\Gamma_{a,inc,fl}$) and $\Gamma_{i \rightarrow a+mol,k}$ (with source strength $\Gamma_{i,inc}$), which represent incident fluid neutral atoms and incident ions, respectively, being reflected and recycled as kinetic atoms and molecules, as determined by the TRIM tables. Adding a separate molecular surface source is hence not needed in this model variant. The fluid neutral atoms now originate exclusively from volume recombination and kinetic-fluid condensation, without which this model would revert to the standard MC approach.

Model HYB-3: automated selection of fluid or kinetic relaunch This model combines those explained above. Based on the local collisionality at each boundary face, it is decided if a particle is recycled kinetically or as a fluid. We calculate the Knudsen number at each grid boundary face of the targets (Kn_{BC}) and introduce two user-defined Knudsen numbers, $\text{Kn}_{\text{BC}}^{\text{l}}$ (lower) and $\text{Kn}_{\text{BC}}^{\text{u}}$ (upper) with $\text{Kn}_{\text{BC}}^{\text{u}} \geq \text{Kn}_{\text{BC}}^{\text{l}}$. When $\text{Kn}_{\text{BC}} \leq \text{Kn}_{\text{BC}}^{\text{l}}$, the fluid approximation for the neutral atoms is deemed sufficiently accurate and model HYB-1 is applied. At boundary faces where $\text{Kn}_{\text{BC}} > \text{Kn}_{\text{BC}}^{\text{u}}$, model HYB-2 is applied. A linear interpolation is applied between $\text{Kn}_{\text{BC}}^{\text{l}}$ and $\text{Kn}_{\text{BC}}^{\text{u}}$ to avoid discontinuities. In principle the user can choose the values for $\text{Kn}_{\text{BC}}^{\text{l}}$ and $\text{Kn}_{\text{BC}}^{\text{u}}$, but in this paper we keep them fixed at $\text{Kn}_{\text{BC}}^{\text{l}} = 0.1$ and $\text{Kn}_{\text{BC}}^{\text{u}} = 0.2$. We formalize this approach by introducing the fluid fraction coefficient R_{fl} , leading to:

$$\Gamma_{a,fl} = R_{fl}(R_{\text{fast},i}\Gamma_{i,\text{inc}} + R_{\text{fast},a}\Gamma_{a,\text{inc},fl}) \quad (6)$$

$$\Gamma_{\text{mol},k} = R_{fl}((1 - R_{\text{fast},i})\Gamma_{i,\text{inc}} + (1 - R_{\text{fast},a})\Gamma_{a,\text{inc},fl}) \quad (7)$$

$$\Gamma_{a,fl \rightarrow a+\text{mol},k} = (1 - R_{fl})\Gamma_{a,\text{inc},fl} \quad (8)$$

$$\Gamma_{i \rightarrow a+\text{mol},k} = (1 - R_{fl})\Gamma_{i,\text{inc}} \quad (9)$$

with

$$R_{fl} = \begin{cases} 1, & \text{if } \text{Kn}_{\text{BC}} \leq \text{Kn}_{\text{BC}}^{\text{l}} \\ 0, & \text{if } \text{Kn}_{\text{BC}} \geq \text{Kn}_{\text{BC}}^{\text{u}} \\ \frac{\text{Kn}_{\text{BC}}^{\text{u}} - \text{Kn}_{\text{BC}}}{\text{Kn}_{\text{BC}}^{\text{u}} - \text{Kn}_{\text{BC}}^{\text{l}}} & \text{otherwise.} \end{cases} \quad (10)$$

The fraction R_{fl} is treated as a fluid, whereas the fraction $1 - R_{fl}$ is still launched kinetically. In EIRENE, the user typically specifies the number of particles that have to be simulated for each source (e.g. 20,000 for the inner target, 2,000 for volume recombination, etc.). For a single target boundary in the HYB-3 model, the user now has to specify a number of particles for $\Gamma_{\text{mol},k}$, $\Gamma_{a,fl \rightarrow a+\text{mol},k}$ and $\Gamma_{i \rightarrow a+\text{mol},k}$. In for example the limit case where $\text{Kn}_{\text{BC}} \geq \text{Kn}_{\text{BC}}^{\text{u}}$ everywhere along the target, this would then mean that CPU time is wasted by simulating all particles for $\Gamma_{\text{mol},k}$ even though they would have no weight. To increase the computational efficiency, the number of particles launched for the sources of $\Gamma_{i \rightarrow a+\text{mol},k}$, $\Gamma_{a,fl \rightarrow a+\text{mol},k}$ and $\Gamma_{\text{mol},k}$ is now automatically scaled according to the relative source strengths. These scaling factors, a_1 , $a_{2,i}$ and $a_{2,a}$, are determined as

$$a_1 = \frac{\sum_{j=1}^N R_{fl,j}(\Gamma_{a,\text{inc},fl,j} + \Gamma_{i,\text{inc},j})A_j}{\sum_{j=1}^N (\Gamma_{a,\text{inc},fl,j} + \Gamma_{i,\text{inc},j})A_j}, \quad (11)$$

$$a_{2,i} = \frac{\sum_{j=1}^N (1 - R_{fl,j})(\Gamma_{i,\text{inc},j})A_j}{\sum_{j=1}^N (\Gamma_{i,\text{inc},j})A_j}, \quad (12)$$

and

$$a_{2,a} = \frac{\sum_{j=1}^N (1 - R_{fl,j})(\Gamma_{a,\text{inc},fl,j})A_j}{\sum_{j=1}^N (\Gamma_{a,\text{inc},fl,j})A_j}, \quad (13)$$

where $j = 1 \dots N$ denotes the summation over all cell faces of a certain target boundary, and A_j denotes the area of a cell face.

2.5 Interpretation as a first-flight model

Since $\text{Kn} = 0.1$ is generally accepted as the transition point between fluid or kinetic treatment, it seems inappropriate to use values for Kn^{t} which are much larger. However, both in Ref. [13] and [14] acceptable accuracy was observed using $\text{Kn}^{\text{t}} = 5$ and $\text{Kn}^{\text{t}} = 10$, respectively. When Kn^{t} is high, some particles will undergo only a few collisions before condensing. These first collisions will, however, often be the ones where the velocity distribution of the neutral atoms still deviates most strongly from the Maxwellian plasma population, providing first-flight information which is hard to capture with a fully fluid neutral model. This first-flight information is obtained for atoms from dissociating molecules and atoms returning to the plasma from the voids for all hybrid variants. For HYB-2 and the kinetic part of HYB-3 ($R_{fl} < 1$),

the first-flight interpretation also extends to the atoms from fast recycling and fast reflection. In this work, we explore the transition point on a broad scale ($\text{Kn}^t = 0.1, 1, 10$ and 100).

Model	Atoms from volumetric recombination	Atoms at plasma-void interfaces	Molecules	Kinetic-fluid condensation	Atoms from fast recycling and reflection at targets	Atoms from dissociating molecules
MC	Kinetic	Kinetic	Kinetic	No	Kinetic	Kinetic
HYB-1	Fluid	Kinetic	Kinetic	Yes	Fluid	Kinetic
HYB-2	Fluid	Kinetic	Kinetic	Yes	Kinetic	Kinetic
HYB-3	Fluid	Kinetic	Kinetic	Yes	Fluid and/or Kinetic depending locally on R_{fl}	Kinetic
Hybrid 1 [8]	Fluid	Kinetic	Kinetic	No	Fluid	Kinetic
Hybrid 2 [8]	Fluid	Kinetic	Kinetic	No	Fluid	Fluid

Table 1: Summary of the fluid-kinetic split-up in the different hybrid approaches.

3 Application to ITER case

3.1 Case set-up

We model a D-only low-power ITER case. The geometry is illustrated in Figure 1. The FV-grid consists of 72 cells in the radial direction and 180 cells in the poloidal direction. Towards the targets, cells are refined to a width of 1-2 mm, to ensure a low discretization error in the region of the strongest plasma-neutral interaction.

The different hybrid neutral models that will be compared in this study are coupled to the same plasma model. Parallel transport is based on Braginskii, but flux-limited. We use an anomalous particle transport coefficient $D^n = 0.3 \text{ m}^2\text{s}^{-1}$, ion/electron heat conductivity coefficients $\chi_i = \chi_e = 1.0 \text{ m}^2\text{s}^{-1}$ and an ion viscocity coefficient $\eta_i = 0.2 \text{ m}^2\text{s}^{-1}$. Drifts and currents are not modeled and the potential equation is not solved. At the core, 20MW of power enters the domain, equally distributed among ions and electrons, and a core particle flux of $9.1 \cdot 10^{21} \text{ s}^{-1}$ represents pellet fueling. At the radial plasma boundaries, decay lengths of 3 cm are imposed for T_i , T_e and n_i . At the puffing location, a flux of $0.5 \cdot 10^{23} \text{ s}^{-1}$ of low-energy D_2 molecules is injected. This puff is launched kinetically in all hybrid methods. At the pumping surface, an absorption coefficient of 0.72% is applied. The targets are made of tungsten, the main chamber wall is made of beryllium and below the dome we use iron for the TRIM tables.

The atomic and molecular reactions which are included in this study are shown in Table 2, and correspond to the complete setup typically used for SOLPS-ITER plasma-edge simulations, but excluding neutral-neutral collisions. For the molecular ions, only a local equilibrium concentration is calculated and their transport is not modeled, as is the standard in present-day EIRENE.

Finally, a choice for L is needed. For simplicity, we take a spatially constant $L = 0.1 \text{ m}$, representing a rough average of the gradient length scales encountered in the problem, as explained in Ref. [15]. It would be physically more appropriate to introduce a spatially dependent $L(\mathbf{r})$, based on local gradient length scales of the plasma, as was done in Ref. [13]. However, these plasma gradient length scales can vary abruptly from timestep to timestep, especially in the presence of statistical noise, which is why we opt for a constant value of L in this work. Assessing whether a more accurate description of $L(\mathbf{r})$ can significantly improve the trade-off between speed-up and accuracy, without introducing convergence issues, is subject of future work.

3.2 Accuracy of the hybrid methods

We assess the accuracy of the hybrid models by studying the obtained plasma profiles along the outer target plate in Figure 2. The combination of three hybrid options for the target boundary conditions,

Reaction	Type	AMJUEL
$D^+ + e \rightarrow D + \text{photon}$	Radiative recombination	2.1.8
$D^+ + 2e \rightarrow D + e$	Three-body recombination	2.1.8
$D + e \rightarrow D^+ + 2e$	Ionization	2.1.5
$D + D^+ \rightarrow D^+ + D$	Charge exchange	3.1.8
$D_2 + e \rightarrow D_2^+ + 2e$	Ionization	2.2.9
$D_2 + e \rightarrow 2D + e$	Dissociation	2.2.5g
$D_2 + e \rightarrow D + D^+ + 2e$	Dissociation	2.2.10
$D_2 + D^+ \rightarrow D_2 + D^+$	Elastic scattering	0.3T
$D_2 + D^+ \rightarrow D_2^+ + D$	Ion conversion	3.2.3
$D_2^+ + e \rightarrow D + D^+ + e$	Dissociation	2.2.12
$D_2^+ + e \rightarrow 2D^+ + 2e$	Dissociative ionization	2.2.11
$D_2^+ + e \rightarrow 2D$	Dissociative recombination	2.2.14

Table 2: Overview of the different reactions included in this study, with the corresponding reaction number in the AMJUEL database [11].

and four values of Kn^t , leads to 12 model variants. Even though all 12 model variants are compared to the same standard MC model, we split the results in three rows (repeating the MC profile) to keep the figures readable. In general, quantitative agreement is achieved between the standard MC approach and most hybrid variants. For HYB-1, the target plasma density, ion particle flux, and ion temperature in the vicinity of the strikepoint show excellent agreement. However, significant errors are made for the ion temperature outward from the strikepoint. This is because the fluid approximation for the neutral atoms is not valid in those regions (low densities and high temperatures result in a large mean-free path). This problem does not occur with HYB-2, as all particles are launched kinetically here (except neutral atoms from volume recombination). The disadvantage of HYB-2 is that it may be inefficient in the vicinity of the separatrix, because particles are relaunched kinetically even though the results from HYB-1 clearly show that the fluid neutral boundary conditions are valid there. The results for HYB-3 show that the benefits of HYB-1 and HYB-2 can indeed be seamlessly combined, allowing to use the fluid neutral boundary conditions where they are valid, while avoiding the large errors on the ion temperature from HYB-1.

Concerning Kn^t , we again find that using $\text{Kn}^t = 1$ or 10 still leads to accurate results, even though $\text{Kn}^t = 0.1$ is regarded as the physical transition threshold. This supports the first-flight argument made in Section 2.5. However, for the extreme value of $\text{Kn}^t = 100$, the agreement clearly starts to deteriorate. The cases with $\text{Kn}^t = 100$ are not further considered due to the substantially poorer accuracy. The cases with $\text{Kn}^t = 0.1$ are accurate, but will only provide a limited speed-up, because a large presence of kinetic neutral atoms remains, as will be shown in the following sections. Hence, to quantify the accuracy, we focus on the cases with $\text{Kn}^t = 1$ and 10. For the peak target plasma density, and peak target ion flux density, HYB-1, HYB-2 and HYB-3 perform similarly, with relative errors of 5-20% for $\text{Kn}^t = 1$, and 10-30% for $\text{Kn}^t = 10$.

3.3 Elimination of the dominant CX reaction from EIRENE

In this section, we show that the dominant CX reaction in the plasma volume can indeed be removed from the EIRENE MC simulation, which is the main motivation of the proposed hybrid methods. Figure 3 shows the molecular density (a), the neutral atom density (b) and the CX mean-free path (c) of the MC reference solution in the divertor. Subsequently, Figure 4 shows the fluid fraction $n_{a,fl}/(n_{a,fl} + n_{a,k})$, i.e. the ratio of the density of neutral atoms represented by the fluid description to the total atom density. The results for $\text{Kn}^t=100$ are omitted due to the poor accuracy. The high fluid fraction towards the core region is an artifact of the minimum fluid density used in SOLPS-ITER (for numerical stability), and

Figure 2: Plasma density (left), ion temperature (middle) and ion particle flux density (right) as a function of the strike-point distance ($s - s_{\text{sep}}$) along the outer target, obtained by model HYB-1 (top row), HYB-2 (middle row) and HYB-3 (bottom row). The standard MC approach (thick blue line) is compared to the hybrid models with $\text{Kn}^t = 0.1$ (black full line), $\text{Kn}^t = 1$ (red dashed-dotted line), $\text{Kn}^t = 10$ (magenta dashed line) or $\text{Kn}^t = 100$ (green dotted line).

Figure 3: Density of neutral molecules (a), neutral atoms (b), and the average charge-exchange mean-free path for the neutral atoms (c) for the MC reference case.

has no physical meaning. Instead, one should focus on the fluid fraction in the divertor legs, where the neutral atom density is high.

From the combination of Figure 3 and Figure 4, it can be seen that the hybrid methods result in a fluid treatment in the regions where a large number of neutral atoms is undergoing a large number of CX collisions (small λ_{CX}), at least for $\text{Kn}^t = 1$ and 10. Even though all fast recycled and fast reflected atoms are treated as fluid in HYB-1, the low fluid fraction in the vicinity of the strikepoint for $\text{Kn}^t = 0.1$ indicates a large contribution of atoms from dissociating molecules. From Figure 3(a) it is also clear that there is a large presence of neutral molecules in the divertor volume. Hence, the speed-up offered by the hybrid methods will be limited by the fraction of the EIRENE CPU time which is used for the kinetic molecules.

3.4 Speed-up and noise reduction

Determining the speed-up of the EIRENE simulation obtained with the hybrid methods is non-trivial. An honest assessment of the speed-up requires comparing simulation times for an equal numerical error [4]. This is complicated by the fact that a given number of MC particles will lead to a different error in the various hybrid approaches. First, the distribution of the launched particles must be chosen differently in the various hybrid methods, because new surface sources are introduced (separate molecular sources and incident fluid neutral atoms which have to be converted to kinetic neutral atoms), and other sources are no longer needed (e.g. volume recombination). Second, both the magnitude and the spatial distribution of the statistical noise will differ between the standard MC approach and each of the 12 hybrid model variants. In this paper, the speed-up is first computed for a sensible choice of the particle distribution, and subsequently corrected according to theoretical error scaling laws [4] to account for the noise reduction offered by the hybrid methods.

The particle distribution for the standard MC case is based on the relative source strengths, but rounded, in-out symmetric, and with a minimum of 1000 particles per source. For the hybrid methods, we propose a sensible choice of the particle distribution, representative of what a code user might select. Table 3 lists the selected particle distributions. The left half of Table 4 shows the resulting EIRENE speed-up ($\text{CPU}_{\text{MC}}/\text{CPU}_{\text{HYB}}$) of the different hybrid variants. To gain further insight in the CPU time taken up by atoms/molecules, we also include results for method HYB-2, but assuming all molecules immediately dissociate at the target, leading to an atom-only (AO) variant. The EIRENE CPU time of this HYB-2 AO model is compared to that of the atom-only version of the fully MC method. The speed-ups are mostly in the range of 3-5 for the cases with molecules. The speed-up for the atom-only hybrid simulations is much larger than for the cases with molecules. This proves that the kinetic simulation of the molecules is limiting the obtained speed-up for the hybrid methods.

Combining Figure 2 and Table 4, we conclude that $\text{Kn}^t = 1$ -10 provides the most promising trade-off between speed-up and accuracy.

Next, we correct the speed-up for the difference in statistical noise between the fully kinetic and hybrid methods. It is expected that the partially fluid approach in the hybrid methods will lead to reduced statistical noise. This is not a certainty though, as the condensation process also introduces

Figure 4: Fluid fraction of the neutral atoms for model variant HYB-1 (top row), HYB-2 (middle row) and HYB-3 (bottom row), using either $\text{Kn}^t = 0.1$ (left column), $\text{Kn}^t = 1$ (middle column) or $\text{Kn}^t = 10$ (right column).

Source	Type	MC	HYB-1	HYB-2	HYB-3
Inner target	$\Gamma_{i \rightarrow a + \text{mol}, k}$	12000	/	12000	$a_{2,i}$ 12000
	$\Gamma_{a, fl \rightarrow a + \text{mol}, k}$	/	/	12000	$a_{2,a}$ 12000
	$\Gamma_{\text{mol}, k}$	/	10000	/	a_1 10000
Outer target	$\Gamma_{i \rightarrow a + \text{mol}, k}$	12000	/	12000	$a_{2,i}$ 12000
	$\Gamma_{a, fl \rightarrow a + \text{mol}, k}$	/	/	12000	$a_{2,a}$ 12000
	$\Gamma_{\text{mol}, k}$	/	10000	/	a_1 10000
Inner PFR	$\Gamma_{i \rightarrow a + \text{mol}, k}$	1000	1000	1000	1000
	$\Gamma_{a, fl \rightarrow a + \text{mol}, k}$	/	1000	1000	1000
Outer PFR	$\Gamma_{i \rightarrow a + \text{mol}, k}$	1000	1000	1000	1000
	$\Gamma_{a, fl \rightarrow a + \text{mol}, k}$	/	1000	1000	1000
Main chamber	$\Gamma_{i \rightarrow a + \text{mol}, k}$	1000	1000	1000	1000
	$\Gamma_{a, fl \rightarrow a + \text{mol}, k}$	/	1000	1000	1000
Vol. rec.	$\Gamma_{i \rightarrow a, k}$	20000	/	/	/
Gas puff	$\Gamma_{\text{mol}, k}$	1000	1000	1000	1000

Table 3: EIRENE particle distribution used for the MC and hybrid simulations. The numbers of particles for the target boundaries for HYB-3 are automatically rescaled with the factors a_1 (Eq. 11), $a_{2,i}$ (Eq. 12) and $a_{2,a}$ (Eq. 13).

	CPUt _{MC} /CPUt _{HYB}				CPUt _{MC} /CPUt _{HYB,corr}			
	Kn ^t = 0.1	Kn ^t = 1	Kn ^t = 10	Kn ^t = 100	Kn ^t = 0.1	Kn ^t = 1	Kn ^t = 10	Kn ^t = 100
HYB-1	1.84	3.19	3.77	3.75	2.96	6.39	14.34	17.28
HYB-2	0.79	2.9	5.66	5.75	1.59	5.48	17.73	20.15
HYB-3	1.69	3.05	3.85	4.08	2.79	6.11	16.24	18.56
HYB-2 AO	0.86	12.49	73.8	127.43				

Table 4: Speed-up of the EIRENE part of the hybrid simulations compared to the fully kinetic approach, including also the atom-only (AO) variant of HYB-2. Left half: speed-up for the particle distributions from Table 3. Right half: speed-up including the correction for the statistical noise, determined by Eq. 14.

statistical noise on the fluid neutral equations. To quantify the statistical noise, we calculate the standard deviation on the plasma density, in the statistically steady-state phase of the simulation, and average it over all cells in the outer divertor leg (*odl*), $\overline{\sigma_{n_i}^{odl}}$. In MC simulations, it is well known that the standard deviation decreases with $1/\sqrt{P}$, where P is the number of particles. This means that the number of particles in the hybrid methods could have been reduced by the square of the ratio of the original average standard deviations, to obtain a similar average standard deviation as the fully kinetic approach, and this reduction in particles linearly translates to the speed-up, resulting in the noise-corrected hybrid CPU time (CPUt_{HYB,corr}):

$$\text{CPUt}_{\text{HYB,corr}} = \text{CPUt}_{\text{HYB}} \cdot \left(\frac{\overline{\sigma_{n_i}^{odl, \text{HYB}}}}{\overline{\sigma_{n_i}^{odl, \text{MC}}}} \right)^2. \quad (14)$$

For the atom-only cases, only the EIRENE speed-up was determined, but not the coupled plasma-neutral solution, hence the noise-corrected speed-up is not given. The noise-corrected speed-up for the cases with molecules is displayed in the right half of Table 4. The speed-up increases to a factor ± 6 for $\text{Kn}^t = 1$ and to 14-18 for $\text{Kn}^t = 10$. This implies that much fewer particles than given in Table 3 could have been used for the hybrid methods. There is little difference between the use of HYB-1, HYB-2, HYB-3. This means that all methods remain interesting candidates for further research. In particular HYB-2 and HYB-3 are promising because they provided an accurate estimate for all physical variables in the computational domain, including the temperatures towards the baffle regions (see Figure 2).

Finally, we note that a non-negligible bias error was observed for the hybrid fluid-kinetic methods. A bias error is a deterministic error originating from the coupling of a statistical model (EIRENE) with a non-linear model (the plasma and fluid neutral equations). While this error is known to be small for the standard MC approach, it was observed that the bias error is much larger for the current implementation of the hybrid methods. Performing the hybrid simulations with more particles further reduced the discrepancy with the fully kinetic model, while decreasing the number of particles increased the discrepancies. The reason for this significant bias error is not yet understood and should be investigated in future research. This also means that it is currently not yet possible to fully exploit the reduction in statistical noise offered by the hybrid methods, because reducing the number of particles increases the bias error.

4 Conclusions and future work

The hybrid fluid-kinetic models for plasma-edge neutrals described in this paper offer an attractive trade-off between accuracy and the ability to remove the dominant CX reaction from the EIRENE MC simulation. The addition of the kinetic-fluid condensation process efficiently deals with kinetic neutrals from plasma-void interfaces, dissociating molecules, and kinetic target boundary conditions, allowing to capture first-flight effects while preventing the explicit simulation of a high number of charge-exchange collisions in the MC approach in high-collisional regions. Using a kinetic-fluid transition threshold Kn^t of 1-10 results in a high fluid fraction in the divertor legs, resulting in a 10-30% loss in accuracy and a factor 5-18 reduction in CPU time for the kinetic neutral solver EIRENE. Using a very high value of $\text{Kn}^t = 100$

was found to be undesirable, because it led to substantially poorer accuracy, but limited additional speed-up. Using fluid neutral boundary conditions everywhere along the target (HYB-1) should be avoided when an accurate prediction of the temperatures towards the baffle regions is desired. This can be solved by relaunching all particles kinetically at the targets (HYB-2) or by using automated optimal selection of the target boundary conditions (HYB-3).

Regarding future work, the methods should be tested on a variety of cases in different tokamaks and different regimes. Finally, the compatibility of the hybrid neutral models with more complex plasma models, such as those containing a self-consistent calculation of drifts and currents, should be investigated. Furthermore, research is needed to understand why the bias error for the proposed hybrid methods is significantly larger than that of the fully kinetic approach, and how this can be resolved.

Acknowledgments

The first author is funded by a PhD fellowship of the Research Foundation - Flanders. This work has been carried out within the framework of the EUROfusion Consortium and has received funding from the Euratom research and training programme 2014-2018 and 2019-2020 under grant agreement No 633053. The views and opinions expressed herein do not necessarily reflect those of the European Commission. The computational resources and services used in this work were provided by the VSC (Flemish Supercomputer Center), funded by the Research Foundation Flanders (FWO) and the Flemish Government – department EWI.

Data Availability Statement

The data that support the findings of this study are available from the corresponding author upon reasonable request.

References

- [1] M. Blommaert et al. “A spatially hybrid fluid-kinetic neutral model for SOLPS-ITER plasma edge simulations”. In: *Nuclear Materials and Energy* 19 (2019), pp. 28–33.
- [2] D. Borodin. “FLUID, KINETIC AND HYBRID APPROACHES FOR EDGE TRANSPORT MODELLING IN FUSION DEVICES”. In: *Nuclear Fusion, Under review* (2021).
- [3] W. Dekeyser et al. “Plasma edge simulations including realistic wall geometry with SOLPS-ITER”. In: *Nuclear Materials and Energy* 27 (2021), p. 100999.
- [4] K. Ghooos et al. “Accuracy and Convergence of Coupled Finite-Volume/Monte-Carlo Codes for Plasma Edge Simulations”. In: *Contributions to plasma physics* 56.6-8 (2016), pp. 616–621.
- [5] A Holm, TD Rognlien, and WH Meyer. “Implementation and assessment of an extended hydrogenic molecular model in UEDGE”. In: *Contributions to Plasma Physics* 60.5-6 (2020), e201900150.
- [6] N Horsten et al. “A hybrid fluid–kinetic neutral model based on a micro–macro decomposition in the SOLPS-ITER plasma edge code suite”. In: *Contributions to Plasma Physics* 60.5-6 (2020), e201900132.
- [7] N. Horsten, G. Samaey, and M Baelmans. “Development and assessment of 2D fluid neutral models that include atomic databases and a microscopic reflection model”. In: *Nuclear Fusion* 57.11 (2017), p. 116043.
- [8] N. Horsten et al. “Application of spatially hybrid fluid–kinetic neutral model on JET L-mode plasmas”. In: *Nuclear Materials and Energy* 27 (2021), p. 100969.
- [9] CFF Karney, DP Stotler, and BJ Braams. “Modeling of neutral plasma in a divertor in the fluid-kinetic transition”. In: *Contributions to Plasma Physics* 38.1-2 (1998), pp. 319–324.

- [10] B. Mortier, M. Baelmans, and G. Samaey. “Kinetic-diffusion asymptotic-preserving Monte Carlo algorithm for Boltzmann-BGK in the diffusive scaling”. In: *arXiv preprint arXiv:2012.08985* (2020).
- [11] D Reiter et al. “The data file AMJUEL: Additional atomic and molecular data for EIRENE”. In: *Forschungszentrum Juelich GmbH* 52425 (2000).
- [12] D. Reiter, M. Baelmans, and P. Börner. “The EIRENE and B2-EIRENE codes”. In: *Fusion science and technology* 47.2 (2005), pp. 172–186.
- [13] M Valentinuzzi et al. “Two-phases hybrid model for neutrals”. In: *Nuclear Materials and Energy* 18 (2019), pp. 41–45.
- [14] W. Van Uytven et al. “A hybrid fluid-kinetic model for plasma-edge neutrals using kinetic-fluid condensation and its application to ITER cases”. In: *47th EPS Conference on Plasma Physics* 45A (2021).
- [15] W. Van Uytven et al. “Assessment of advanced fluid neutral models for the neutral atoms in the plasma edge and application in ITER geometry”. In: *In preparation* ().
- [16] W. Van Uytven et al. “Implementation of a separate fluid-neutral energy equation in SOLPS-ITER and its impact on the validity range of advanced fluid-neutral models”. In: *Contributions to Plasma Physics* 60.5-6 (2020), e201900147.
- [17] S. Wiesen et al. “The new SOLPS-ITER code package”. In: *Journal of nuclear materials* 463 (2015), pp. 480–484.