

Effect of time step, neutral–neutral collisions, and an underrelaxation scheme on the numerical convergence of SOLPS-ITER plasma boundary simulations with kinetic neutrals

Original

Effect of time step, neutral–neutral collisions, and an underrelaxation scheme on the numerical convergence of SOLPS-ITER plasma boundary simulations with kinetic neutrals / Van Uytven, W.; Subba, F.; Wiesen, S.; Horsten, N.; Tang, Z.; Dekeyser, W.. - In: PHYSICS OF PLASMAS. - ISSN 1070-664X. - 32:10(2025). [10.1063/5.0292725]

Availability:

This version is available at: 11583/3008419 since: 2026-03-09T12:08:05Z

Publisher:

American Institute of Physics

Published

DOI:10.1063/5.0292725

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

RESEARCH ARTICLE | OCTOBER 28 2025

Effect of time step, neutral–neutral collisions, and an underrelaxation scheme on the numerical convergence of SOLPS-ITER plasma boundary simulations with kinetic neutrals

W. Van Uytven ; F. Subba ; S. Wiesen ; N. Horsten ; Z. Tang ; W. Dekeyser 



Phys. Plasmas 32, 103907 (2025)

<https://doi.org/10.1063/5.0292725>



Articles You May Be Interested In

Numerical implications of including drifts in SOLPS-ITER simulations of EAST

Phys. Plasmas (February 2024)

Experimental investigation and SOLPS-ITER modeling of Ne-seeded radiative divertor H-modes plasma on EAST

Phys. Plasmas (May 2019)

Self-consistent simulation of transport and turbulence in tokamak edge plasma by coupling SOLPS-ITER and BOUT++

Phys. Plasmas (January 2019)

09 March 2026 11:53:42



AIP Advances

Why Publish With Us?

-  **21DAYS**
average time to 1st decision
-  **OVER 4 MILLION**
views in the last year
-  **INCLUSIVE**
scope

[Learn More](#)



Effect of time step, neutral-neutral collisions, and an underrelaxation scheme on the numerical convergence of SOLPS-ITER plasma boundary simulations with kinetic neutrals

Cite as: Phys. Plasmas **32**, 103907 (2025); doi: 10.1063/5.0292725

Submitted: 23 July 2025 · Accepted: 9 October 2025 ·

Published Online: 28 October 2025



View Online



Export Citation



CrossMark

W. Van Uytven,^{1,a)} F. Subba,² S. Wiesen,³ N. Horsten,¹ Z. Tang,⁴ and W. Dekeyser¹

AFFILIATIONS

¹KU Leuven, Department of Mechanical Engineering, Celestijnenlaan 300A, 3001 Leuven, Belgium

²NEMO Group, Politecnico di Torino, Torino, Italy

³DIFFER-Dutch Institute for Fundamental Energy Research, De Zaale 20, 5612 AJ Eindhoven, The Netherlands

⁴KU Leuven, Department of Computer Science, Celestijnenlaan 200A, 3001 Leuven, Belgium

^{a)} Author to whom correspondence should be addressed: wim.vanuytven@kuleuven.be

ABSTRACT

The numerical convergence of plasma boundary simulations with kinetic neutrals in SOLPS-ITER is studied by performing a cross-parameter scan on the time step Δt and the number of Monte Carlo particles per iteration P . A low-power D-only DEMO case, a high-power D + He + Ar DEMO case and a D + He artificial slab case are used to perform the analysis. The numerical scans reveal that for a given time step, the bias error scales with $1/P$, as expected from theory. For a given P , the bias error scales with Δt^n , with $n \in [1, 2]$. It is found that the D + He + Ar DEMO case requires a large amount of impurity Monte Carlo particles per iteration to achieve acceptable numerical errors. Furthermore, it is explained that neutral-neutral collisions introduce an additional bias contribution, but that it is negligible compared to the bias due to the plasma-neutral coupling for the studied cases. Additionally, it is shown that global particle imbalances are linked to the numerical convergence. Finally, it is shown that an underrelaxation scheme for the kinetic neutral source terms is able to reduce the bias errors by one to two orders of magnitude.

© 2025 Author(s). All article content, except where otherwise noted, is licensed under a Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>). <https://doi.org/10.1063/5.0292725>

I. INTRODUCTION

Plasma boundary codes such as SOLPS-ITER^{1,2} are used to study the power exhaust in magnetically confined fusion devices. The two main physics modules of SOLPS-ITER are B2.5 and EIRENE. B2.5 is a finite-volume (FV) fluid code that solves transport equations for the plasma.³ EIRENE is a kinetic particle-tracing Monte Carlo (MC) code, which simulates the neutral particles in the plasma edge.⁴

This paper investigates the numerical errors arising from the coupling of the stochastic MC neutrals with the FV plasma model, henceforth abbreviated as the FV-MC coupling. Because plasma boundary simulations can be very demanding in terms of computational cost, it is of great interest to know how one can achieve a desired numerical accuracy with the least amount of computational resources. The trade-off between numerical accuracy and computational cost has already been studied and optimized intensively in the work of Ghooes *et al.*^{5–8} However,

important questions remain that were not yet answered there, but that are relevant for realistic simulations of, for example, ITER and (EU-)DEMO. All 1D and 2D simulations performed in the aforementioned works^{5–8} were deuterium-only cases without neutral-neutral collisions. Furthermore, the vast majority of the analyses were performed with a fixed time step for the plasma model. Hence, three unanswered questions are: (i) What is the effect of neutral-neutral collisions on the error scaling? (ii) What is the effect of the time step of the plasma solver on the error scaling? (iii) What is the effect of impurity species on the error scaling? Additionally, we will study the performance of an underrelaxation scheme, which can further improve the numerical convergence.

The paper is structured as follows: Sec. II gives an overview of the FV-MC coupling, the different numerical error contributions in FV-MC plasma boundary models with their expected scalings, and the

proposed underrelaxation scheme. Next, Sec. III details the case setups used for the convergence studies and error estimations. We use a (EU-)DEMO geometry with a deuterium-only model and a D + He + Ar model. Furthermore, we employ a D + He artificial slab case, which is less physically relevant but which is convenient for numerical experiments. In Sec. IV, the results of the convergence studies are analyzed and discussed. Finally, Sec. V ends with the conclusions and recommendations for future work.

II. THE FLUID-KINETIC COUPLING IN SOLPS-ITER AND ITS NUMERICAL ERROR CONTRIBUTIONS

A. FV-MC coupling

The B2.5-EIRENE coupling scheme is represented schematically in Fig. 1. B2.5 provides incident fluxes and plasma backgrounds to EIRENE. EIRENE then performs the particle tracing and ultimately provides sources of particles, momentum, and ion/electron energy for the right-hand side of the plasma equations. In a single B2.5 inner iteration, an update is performed for the density and parallel velocity of each ion species, the electron temperature, the combined (over all ion species) ion temperature, and the plasma potential when including currents. This update is performed by using a segregated semi-implicit false-time-stepping approach, with time step Δt and relaxation factor α . It is possible to perform multiple B2.5 inner iterations per EIRENE call, but this option is not considered here, as explained in Sec. II H. When EIRENE includes neutral-neutral collisions, it is also possible to perform inner iterations on EIRENE. This is typically not done in SOLPS-ITER, as explained in Sec. II E.

This paper focuses only on finding the steady-state solution of the FV-MC system; hence, we run EIRENE in steady-state mode, meaning that all neutral particles are followed from birth until death (i.e., ionization, pumping, or crossing the core boundary). Hence, the time step Δt in this paper refers exclusively to the plasma model.

B. Numerical error contributions

The total numerical error on the plasma solution, ϵ_{num} , can be written as⁵

$$\epsilon_{\text{num}} = \epsilon_d + \epsilon_s + \epsilon_b + \epsilon_c. \quad (1)$$

ϵ_d is the discretization error, resulting from the finite resolution of the FV mesh. This error was recently studied for a series of EU-DEMO

cases with a fluid approximation for the neutrals.⁹ There, ϵ_d was found to be in the 4%–25% range for the 96×36 grid that is used here. In this contribution, all simulations are performed on the same 96×36 mesh. We do not study the discretization error here, but expect it to be of the same order as in Ref. 9.

ϵ_s is the statistical error, following directly from the Monte Carlo noise of the kinetic neutral solver EIRENE. ϵ_s has zero mean, and scales with $\frac{1}{\sqrt{P}}$ ⁵ where P is the number of MC particles launched per iteration. ϵ_s can be made arbitrarily small by averaging the plasma solution over sufficient iterations, once a statistical steady state is reached (see Sec. II C).

The finite-sampling bias error, ϵ_b , is the deterministic error which arises due to the coupling of the statistical noise with the non-linear plasma model. Contrary to the statistical error, it does not disappear when time-averaging the plasma solution. The convergence error, ϵ_c , is the remaining imbalance in the discretized conservation equations in the statistical steady state. Similarly as the finite-sampling bias, it does not disappear when averaging the plasma states in the statistical steady state. The finite-sampling bias and the convergence error are often summed as $\epsilon_{\text{bc}} = \epsilon_b + \epsilon_c$, henceforth called the bias error. In the work of Ghoois *et al.*, it was shown that the bias error is theoretically expected to scale with $\frac{1}{P}$ and this was confirmed in each numerical experiment.

C. Random noise averaging strategy

Random Noise Averaging (RNA) emerged from the work of Ghoois *et al.* as the most efficient method to minimize the numerical error for a given computational time.⁸ Random seeds are used for each EIRENE iteration (as opposed to the correlated sampling approach), and the neutral sources are not averaged over the iterations at runtime (as opposed to the Robbins–Monro approach). The simulation is run until a statistical steady state is reached, where all plasma variables and residuals oscillate around constant average values. Then, a post-processing average of this statistical steady state is calculated. The statistical error then reduces as $\frac{\sqrt{T}}{\sqrt{IP}}$, where I is the number of iterations in the post-processing average, and T is a correlation time between successive iterations. Hence, the statistical error can be effectively removed by averaging over sufficient iterations. The remaining bias error was found to be acceptable, unless a very low number of particles per iteration is used. However, RNA may not be the optimal strategy for cases that do suffer from large bias errors.

D. Effect of Δt on the bias

For a given number of particles per iteration, the statistical error and bias error also depend on the time step Δt of the plasma solver. In Ref. 10, it was shown analytically that the statistical error is expected to scale with $\sqrt{\Delta t}$ and the bias with Δt . This was confirmed with both 0D and 1D numerical experiments. A smaller time step leads to reduced statistical noise on the plasma quantities, because in each time step the plasma adapts less to each particular instance of the randomly noisy EIRENE sources, hence better filtering out the noise. In this paper, we will quantify the effect of Δt on the bias error for 2D SOLPS-ITER cases.

E. Effect of neutral-neutral collisions

Presently, it is common to include neutral-neutral collisions (NNCs) in EIRENE simulations, especially if the (sub-)divertor neutral

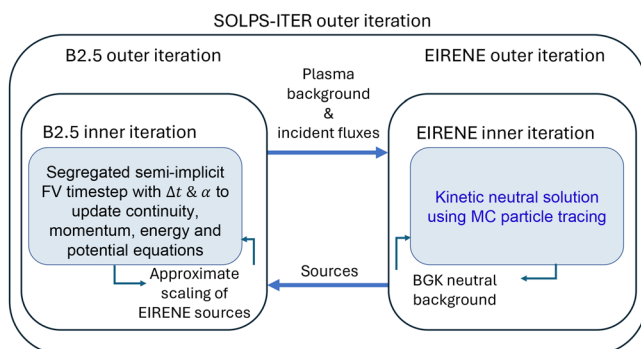


FIG. 1. Schematic representation of the FV-MC coupling scheme between B2.5 and EIRENE in SOLPS-ITER.

densities are expected to be high. Consider, as an example, the steady-state kinetic Boltzmann equation for the atoms with ionization, recombination, charge exchange, and atom–atom elastic collisions

$$\begin{aligned} \mathbf{v} \cdot \nabla f_a(\mathbf{v}) &= f_i(\mathbf{v})n_e K_r - f_a(\mathbf{v})n_e K_i \\ &- f_a(\mathbf{v}) \int_{\mathbf{v}'} \sigma_{\text{CX}}(E_c) \|\mathbf{v} - \mathbf{v}'\| f_i(\mathbf{v}') d\mathbf{v}' \\ &+ f_i(\mathbf{v}) \int_{\mathbf{v}'} \sigma_{\text{CX}}(E_c) \|\mathbf{v} - \mathbf{v}'\| f_a(\mathbf{v}') d\mathbf{v}' \\ &+ \int_{\mathbf{v}'} \int_{\mathbf{v}_1'} \int_{\mathbf{v}_1} (\sigma_{\text{el}}(\mathbf{v}', \mathbf{v}_1'; \mathbf{v}, \mathbf{v}_1) \|\mathbf{v}' - \mathbf{v}_1'\| \dots \\ &\times (f_a(\mathbf{v})f_a(\mathbf{v}_1) - f_a(\mathbf{v}')f_a(\mathbf{v}_1')) d\mathbf{v}' d\mathbf{v}_1' d\mathbf{v}_1). \end{aligned} \quad (2)$$

Here, \mathbf{v} is the particle velocity vector, $f_a(\mathbf{v})$ is the atom velocity distribution function (the spatial dependence is omitted everywhere in the notation), $f_i(\mathbf{v})$ is the ion velocity distribution function assumed to be a drifting Maxwellian distribution based on ion density, velocity, and temperature. n_e is the electron density, n_i is the ion density, K_i is the ionization rate coefficient, K_r is the recombination rate coefficient, σ_{CX} is the microscopic cross section for charge exchange, E_c is the center-of-mass kinetic energy, and σ_{el} is the differential cross section for atom–atom elastic collisions. \mathbf{v}' denotes the pre-collisional particle velocity vector.

Without elastic collisions ($\sigma_{\text{el}} = 0$), the kinetic equation is linear in $f_a(\mathbf{v})$. In practice, this means that EIRENE can find the steady-state solution in a single iteration on a given plasma background. When NNCs are included, Eq. (2) becomes non-linear in $f_a(\mathbf{v})$. NNCs in present-day EIRENE are not treated as collisions between different test particles, as this would destroy the trivially parallel nature of the MC procedure. Instead, EIRENE makes use of the BGK (Bhatnagar–Gross–Krook) approximation.^{11,12}

$$\begin{aligned} \mathbf{v} \cdot \nabla f_a(\mathbf{v}) &= f_i(\mathbf{v})n_e K_r - f_a(\mathbf{v})n_e K_i \\ &- f_a(\mathbf{v}) \int_{\mathbf{v}'} \sigma_{\text{CX}}(E_c) \|\mathbf{v} - \mathbf{v}'\| f_i(\mathbf{v}') d\mathbf{v}' \\ &+ \int_{\mathbf{v}'} \sigma_{\text{CX}}(E_c) \|\mathbf{v} - \mathbf{v}'\| f_i(\mathbf{v}') f_a(\mathbf{v}') d\mathbf{v}' \\ &- \nu_{\text{BGK}} (f_a(\mathbf{v}) - f_{a,\text{BGK}}(\mathbf{v})), \end{aligned} \quad (3)$$

where ν_{BGK} is the collision frequency and $f_{a,\text{BGK}}(\mathbf{v})$ is the BGK neutral background. It is a perfect drifting Maxwellian based on the moments of $f_a(\mathbf{v})$. Equation E is then still non-linear in $f_a(\mathbf{v})$ because $f_{a,\text{BGK}}(\mathbf{v})$ depends on $f_a(\mathbf{v})$. The final step to ensure that EIRENE does not have to consider interaction between different Monte Carlo neutrals during the particle tracing in a given iteration k is to base $f_{a,\text{BGK}}(\mathbf{v})$ at iteration k on $f_a(\mathbf{v})$ of the previous iteration $k - 1$. $f_{a,\text{BGK}}^k(\mathbf{v})$, is given by

$$f_{a,\text{BGK}}^k(\mathbf{v}) = n_a^{k-1} \left(\frac{m_a}{2\pi k_B T_a^{k-1}} \right)^{\frac{3}{2}} \exp \left(- \frac{m_a}{2T_a^{k-1}} \|\mathbf{V}_a^{k-1} - \mathbf{v}\|^2 \right), \quad (4)$$

where m_a is the atom mass and where

$$n_a^{k-1} = \int_{\mathbf{v}} f_a^{k-1}(\mathbf{v}) d\mathbf{v}, \quad (5)$$

$$\mathbf{V}_a^{k-1} = \frac{1}{n_a^{k-1}} \int_{\mathbf{v}} \mathbf{v} f_a^{k-1}(\mathbf{v}) d\mathbf{v}, \quad (6)$$

$$T_a^{k-1} = \frac{2}{3n_a^{k-1}} \int_{\mathbf{v}} \frac{m_a}{2} \|\mathbf{v}\|^2 f_a^{k-1}(\mathbf{v}) d\mathbf{v} - \frac{m_a}{3} \|\mathbf{V}_a^{k-1}\|^2. \quad (7)$$

The non-linearity of the kinetic equation with BGK NNCs can thus be resolved by performing inner EIRENE iterations for a given plasma background, as shown in Fig. 1. However, this is not common practice in SOLPS-ITER. Instead, the plasma and the neutral BGK background are relaxed together in a single loop. The non-linearity of the kinetic equation with BGK NNCs introduces an additional component to the bias error. This bias is also expected to decrease with $\frac{1}{p}$, but from Fig. 1 it is evident that this bias contribution is independent of the time step and relaxation factor used in the plasma model.

F. Effect of impurity species

The way in which the numerical errors scale with respect to P and Δt is not expected to change when going from D-only to multi-species cases. However, the absolute values of the errors can differ. For example, in Ref. 13 it was shown that the global particle imbalance (which is a symptom of non-convergence) for helium was substantially larger than for deuterium.

EIRENE has been optimized throughout the years to efficiently deal with highly collisional neutrals such as charge-exchange-dominated deuterium atoms. However, the Monte Carlo techniques used in EIRENE might not be optimal for absorption-dominated impurity neutrals. Although absorption-dominated particles require less CPU time per particle, they also provide less statistical information per particle, potentially resulting in a poorer ratio of statistical error to CPU time (see, for example, Ref. 14). The details of the Monte Carlo techniques are, however, not the focus of this paper, and we use the default EIRENE settings in this regard.

G. Underrelaxation scheme for the kinetic source terms

In this paper, we study an additional scheme to reduce the noise by averaging the source terms over multiple iterations before inserting them back into the plasma equations. The general idea is that plasma boundary simulations typically use small time steps, such that the neutral sources calculated on the plasma background of some 10 or 100 iterations ago still contain useful information, especially when nearing or in the steady-state phase. The concept of using a form of runtime-averaging of the MC sources to reduce the noise and bias is not novel. A variant of the scheme that will be used here (using the true moving average and resetting this average after N iterations) was described in Ref. 13, and multiple variants were implemented in the code around the time of that report. It is also mentioned there that time-averaging of MC sources is applied in the code SOLDOR/NEUT2D according to Ref. 15. However, there is a lack of rigorous assessment of these schemes in the context of comprehensive convergence analyses.

1. Definition of the underrelaxation scheme

Let S_a^n , $S_a^{m||}$, and S_a^E be the particle, parallel momentum, and ion energy source for each plasma species a and S_e^E the electron energy source. We refer here specifically to the noisy sources due to neutral interaction as calculated directly by EIRENE, before the ion sources are summed over all species and before the conversion from total energy sources to internal heat sources as required by B2.5. For each

$S \in \{S_a^n, S_a^{m||}, S_a^E, S_e^E\}$, we obtain at iteration i a noise-reduced source term \bar{S} by recursively taking

$$\bar{S}_i = \frac{1}{N} S_i + \frac{N-1}{N} \bar{S}_{i-1}, \quad \text{with } N = \min(i, N_{\text{ur}}), \quad (8)$$

where N_{ur} determines the degree of underrelaxation. We will refer to this scheme simply as underrelaxation in the remainder of the paper. It can also be described as a recursive approximation to a moving average. The difference is that for a true moving average, the weights would be equal to $1/N$ for the last N source terms, and zero for older contributions. However, then, all sources $S_a^n, S_a^{m||}, S_a^E,$ and S_e^E have to be saved N times. In the underrelaxation scheme, the contribution from each previous source term is $\frac{1}{N_{\text{ur}}} (\frac{N_{\text{ur}}-1}{N_{\text{ur}}})^j$, where j indicates the number of iterations that have passed. Because of the recursive nature, the required data storage does not scale with N_{ur} .

When N_{ur} is set to the total number of iterations, the scheme becomes the Robbins–Monro approach.¹⁶ In that case, all noise will eventually be removed, but convergence becomes extremely slow.⁵ Indeed, as more iterations pass, the rate at which the averaged sources can still change keeps decreasing. In this paper, we consider the underrelaxation scheme with a fixed N_{ur} , e.g., 10, 100, or 1000. In that case, source terms from many iterations ago ($j > N_{\text{ur}}$) are eventually forgotten, and the rate of change of the averaged source terms remains constant. For the simulations in this paper we only consider the underrelaxation scheme and not other variants such as the true moving average, a moving average that is reset after a certain time, or the Robbins–Monro approach.

2. Correction for improved flux conservation

A correction to the scheme is employed to improve global particle, momentum, and energy conservation. EIRENE employs stratified sampling, i.e., the source terms in the domain are calculated as a sum of contributions from N_{strat} so-called “strata.” The types of strata used here are (i) the recombination of each isonuclear sequence on each solid surface or plasma-void interface; (ii) gas puff sources; (iii) neutral sources due to volumetric recombination. The stratified sampling allows to define for each stratum s the normalized sources $S_{a,s}^n = \frac{S_{a,s}^n}{\Gamma_s^n}$, $S_{a,s}^{m||} = \frac{S_{a,s}^{m||}}{\Gamma_s^{m||}}$, $S_{a,s}^E = \frac{S_{a,s}^E}{\Gamma_s^E}$, and $S_e^E = \frac{S_e^E}{\Gamma_s^E}$, where Γ_s^n is the particle flux strength of stratum s . The final form of the underrelaxation scheme is to apply Eq. (8) to the normalized sources S , and afterward multiply by the stratum-specific particle flux corresponding to the most recent plasma state.

3. Expected behavior

The advantage of the underrelaxation scheme is that it provides noise reduction without any increase in computational cost or memory requirements. The disadvantage of the scheme is that the moving-averaged source terms, which are plugged back into the plasma equations, are lagging behind on the most recent plasma state. This could potentially slow down the convergence, or create instabilities that lead to oscillations or code divergence. The performance of the underrelaxation scheme will be studied here in detail. To the best of the knowledge of the authors, these schemes have fallen out of use. A possible reason is that in Ref. 13 it was shown that the source averaging does

not work well in conjunction with inner B2.5 iterations, and it was chosen to further pursue the path of inner iterations. The authors are not aware of any publication providing a detailed report of the performance of the underrelaxation scheme in the form presented here.

It can also be considered to apply the underrelaxation scheme to the BGK neutral background. However, the results of this paper will indicate that this is not a priority and hence this has not yet been done.

H. Note on internal iterations

A popular approach to speed-up B2.5-EIRENE simulations is to use internal iterations, whereby multiple B2.5 false time steps are performed after a call to EIRENE. In the simplest form, the source terms from EIRENE are kept constant during the inner iterations. This method quickly leads to instabilities and is not used in practice. A first improvement is to linearly scale the stratum-based sources with the flux strengths of the most recent plasma state, identically as explained in Sec. II.G.2. Even with this improvement, the use of internal iterations was found to severely limit the maximum allowable time step.^{13,17} Therefore, advanced internal iteration cycles were developed, with single internal iterations on all equations but multiple (up to 99) internal iterations for only the continuity equations. The details of this advanced iteration scheme are described in chapter 5.2–5.4 of Ref. 13. Despite their proven benefits, we decide not to use internal B2.5 iterations in this work, for three reasons:

First, while (advanced) inner iterations were shown to decrease residuals and improve the global balances,^{13,17} they are expected to increase the finite sampling bias, because at each iteration the plasma solution adapts more to a particular instance of the noisy sources, partially undoing the filtering effect discussed in Sec. II.D. This was also concluded in Ref. 18. This makes it very difficult to estimate the net effect of inner iterations on the different error contributions from a theoretical point of view. Especially the advanced iteration schemes make it difficult to predict how the solution will be precisely affected.

Second, the analysis performed here is already multi-dimensional (time step, number of particles, inclusion/exclusion of NNCs, and underrelaxation parameter). Adding an additional dimension for the internal iterations would make the study untractable.

Third, it was shown that the source averaging does not work well in conjunction with internal B2.5 iterations, as mentioned in Sec. II.G. We choose to focus on the underrelaxation scheme here.

I. Global particle balances

It is common practice in B2.5-EIRENE simulations to study the global (i.e., over the entire computational domain) balances for particles, parallel momentum, and energy of the coupled fluid–kinetic system. In the steady state, the net influx of particles, momentum, or energy should be equal to the sum of the net outfluxes and dissipation terms. If these global balances are not satisfied, it is an indication of the presence of numerical errors. The advantage of tracking the global imbalances is that it does not require a more computationally expensive reference solution. The disadvantage is that it provides no information on the magnitude of local error contributions and no strict guarantee of convergence. Indeed, it is theoretically possible that the global balances would be closely satisfied, while the residuals at the

level of individual cells or regions are not satisfied. However, in practice, it can be assumed that there is some proportional relationship between the global imbalances and the convergence errors.

In this paper we will only focus on the particle balances, assuming that qualitatively similar trends would be observed for momentum and energy. The particle balances were found to be most critical to satisfy in previous work.^{13,17} In this work, the global imbalance of particles is defined as

$$\Delta\Gamma^\alpha = \frac{|\Gamma_{\text{in}}^\alpha - \Gamma_{\text{out}}^\alpha|}{\Gamma_{\text{in}}^\alpha}, \quad (9)$$

where $\Gamma_{\text{in}}^\alpha$ represents the sum of all inputs of particles belonging to isonuclear sequence α and $\Gamma_{\text{out}}^\alpha$ represents the output. $\Gamma_{\text{in}}^\alpha$ consists of gas puff injection and core boundary fluxes. Wall sputtering and pellet fueling are not considered here. $\Gamma_{\text{out}}^\alpha$ is dominated by the pumped neutrals. Minor contributions are neutrals that cross the core boundary and ions that are directly absorbed at a surface. Note that in high-recycling and detached regimes, the net influxes and outfluxes of particles can be one to two orders of magnitude smaller than the target recycling fluxes. This adds to the numerical challenge of satisfying the balances to a high accuracy.

III. SIMULATION SETUP

We perform numerical convergence analyses on three different cases. Two cases employ an EU-DEMO geometry (Fig. 2), taken from Ref. 19. At the time of writing, the EU-DEMO equilibria are being revised. However, the important point is that the numerical tests are performed on a reactor-sized geometry. The first case is deuterium-only while the second contains deuterium, helium, and argon, and has a fusion-relevant core power flux. The third case is an artificial slab geometry containing deuterium and helium. The test cases are described in more detail in the following sections. All cases are run without drifts and currents. The considered atomic and molecular reactions are identical to the ones in Ref. 19.

A. D-only DEMO case

The D-only DEMO case has strongly simplified boundary conditions. We impose a fixed core temperature of $T_i = T_e = 1000$ eV, a fixed core ion density of $5 \times 10^{19} \text{ m}^{-3}$, and a D_2 gas puff of $1 \times 10^{23} \text{ s}^{-1}$. 0.6% of incident neutrals are pumped away at the bottom of the divertor. The anomalous transport is defined through the anomalous diffusion coefficient, $D_i^{(\text{AN})}$, the anomalous viscosity $\eta_i^{(\text{AN})} = \nu_i^{(\text{AN})} n_i m_i$, and the anomalous ion and electron heat conduction $\kappa_i^{(\text{AN})} = \chi_i^{(\text{AN})} n_i$ and $\kappa_e^{(\text{AN})} = \chi_e^{(\text{AN})} n_e$. The anomalous transport coefficients are spatially constant with $D_i^{(\text{AN})} = \chi_{i/e}^{(\text{AN})} = \nu_i^{(\text{AN})} = 1 \text{ m}^2/\text{s}$.

B. D + He + Ar DEMO case

The D + He + Ar DEMO case is based on Ref. 19. It contains deuterium, helium, and argon, with all their charge states. A total heat flux of 150 MW enters the domain from the core, equally distributed among ions and electrons. The core particle flux of deuterium ions is $7.5 \times 10^{21} \text{ s}^{-1}$, and that of He^{2+} ions is $7.1 \times 10^{20} \text{ s}^{-1}$. Zero-flux boundary conditions at the core are imposed for He^+ and for all argon ionization stages. The same D_2 gas puff of $1 \times 10^{23} \text{ s}^{-1}$ is present, together with a $1.5 \times 10^{20} \text{ s}^{-1}$ puff of argon atoms. The locations of

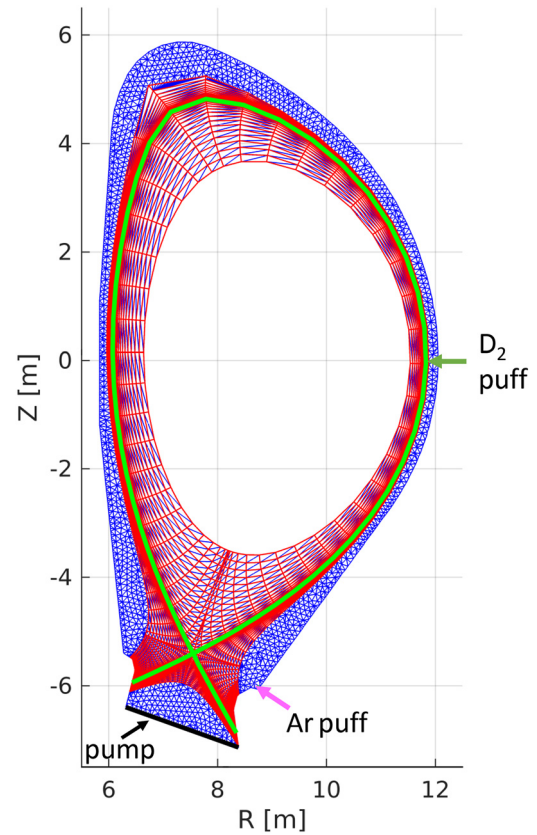


FIG. 2. Computational mesh used for the presented DEMO cases. In red, the B2.5 FV plasma grid. In blue, the triangle mesh for EIRENE. Inside the B2.5 mesh, each quadrilateral is split into two triangles.

the gas puffs are given in Fig. 2. To mimic an H-mode transport barrier, radially dependent transport coefficients are used, identically as the ones in Ref. 9.

C. D + He artificial slab case

Some of the findings in this paper are supported by simulations on an artificial slab case, as shown in Fig. 3. The details of the slab geometry are explained in earlier work.^{20,21} The slab case lends itself excellently to numerical experiments, due to its fast convergence and high mesh quality near the target plates. The fast convergence means that many simulations can be performed at acceptable computational cost and time. The high mesh quality ensures that conclusions about the FV-MC coupling are not contaminated by artifacts arising from poor mesh quality.

The slab case contains deuterium and helium. A fixed core particle flux of $2.5 \times 10^{22} \text{ s}^{-1}$ is imposed for deuterium ions and $2.5 \times 10^{21} \text{ s}^{-1}$ for the He^{2+} ions. 4 MW enters from the core, equally distributed among ions and electrons. The artificial slab case does not have void regions. Each quadrilateral is split in two triangles for EIRENE. The artificial cuts make sure that the core and private flux regions are magnetically separated. At the private flux boundaries, 5% of incident neutrals and ions are pumped away. Because the case is poloidally

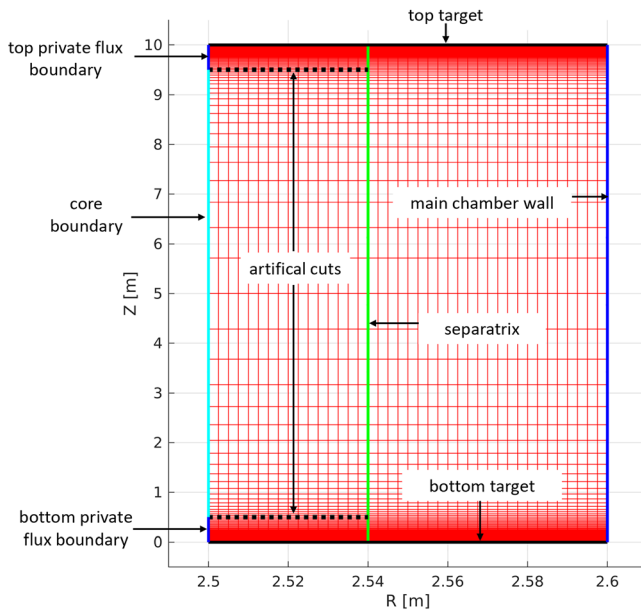


FIG. 3. B2.5 computational mesh for the artificial slab cases.

symmetric, we do not make a distinction between the top and bottom targets in the result discussion. The anomalous transport coefficients are spatially constant with $D_i^{(AN)} = \chi_{i/e}^{(AN)} = 1$ and $\nu_i^{(AN)} = 0.2 \text{ m}^2/\text{s}$.

D. Scan on numerical parameters

To answer the research questions posed in Sec. I, we perform a cross-parameter scan on the time step of the plasma fluid solver, Δt , and the number of particles per iteration, P .

The D-only DEMO case and the artificial slab are scanned with $\Delta t = \{10^{-4}, 10^{-5}, 10^{-6} \text{ s}\}$ and $\{P/100, P/10, P, P \times 10\}$, where, for example, $P/10$ means that the reference particle distribution is reduced by a factor 10 for each stratum. The reference particle distribution for each case is given in Table I. Because the multi-species DEMO cases are most demanding in terms of computational cost and

time, only a subset of the scans is performed for them. The relaxation factor α is kept constant at 0.5. Inside the separatrix, the time step is increased by a factor 1000 for the D-only DEMO cases and by a factor 100 for the D + He + Ar DEMO cases. It was checked that this multiplication has no effect on the numerical errors, because almost no neutral particles reach this region.

Furthermore, the scans for the D-only DEMO case and the artificial slab are performed both with and without neutral-neutral collisions. The multi-species DEMO cases are only run with NNCs included. The simulations are performed in the new unstructured version of SOLPS-ITER,²² but the FV meshes are still structured.

IV. RESULTS

A. D-only DEMO case

1. Effect of neutral-neutral collisions on solution

The physical interpretation of the effect of NNCs is not the focus of this paper. What is important is that the NNCs have a non-negligible effect on the solution; otherwise, there would be no point in analyzing their effect on the numerical convergence. Figure 4 displays the ion temperature and electron density along the inner and outer target with and without NNCs for the numerically most accurate case ($P \times 10$ and $\Delta t = 10^{-6} \text{ s}$). Target quantities in this paper are plotted as a function of $s - s_{\text{sep}}$, where s is the length along the target plate running from private flux to scrape-off layer, and s_{sep} is the length corresponding to the strikepoint location. The NNCs have a distinguished impact on the steady-state solution, causing, for example, a $\sim 50\%$ difference on the peak inner target electron density.

2. Bias contribution due to neutral-neutral collisions

In this section, we study the bias error of the BGK NNC mechanism in an isolated way, by performing inner EIRENE iterations on a fixed plasma background. Figure 5(a) shows the ion heat source term coming from EIRENE, along the bottom-most flux tube of the plasma grid in the inner private flux region. The statistical noise is removed by averaging over I iterations, once the neutral model is in a statistical steady state. For the case with $P \times 10$, we used $I = 100$, and for the cases with less particles I is determined by keeping $P \times I$ constant. Without NNCs, the result is independent of the number of particles per iteration, as expected. With NNCs, we indeed observe a

TABLE I. Reference particle distribution P for the three cases.

	D + He artificial slab		D-only DEMO	D + He + Ar DEMO		
	D	He	D	D	He	Ar
Inner target	2.5×10^4	2.5×10^4	2.5×10^4	2.5×10^4	2.5×10^5	2.5×10^5
Outer target	2.5×10^4	2.5×10^4	2.5×10^4	2.5×10^4	2.5×10^5	2.5×10^5
Inner PF	2×10^4	2×10^4	2×10^3	2×10^3	2×10^4	2×10^4
Outer PF	2×10^4	2×10^4	2×10^3	2×10^3	2×10^4	2×10^4
MCW	2×10^4	2×10^4	2×10^3	2×10^3	2×10^4	2×10^4
Vol. Rec.	1×10^4	1×10^4	7×10^3	7×10^3	7×10^4	7×10^4
Gas puff	1×10^4	1×10^4	...	1.10^5
Tot/species	1.2×10^5	1.2×10^5	7.3×10^4	7.3×10^4	6.3×10^5	7.3×10^5
Total	2.4×10^5		7.3×10^4	1.43×10^6		

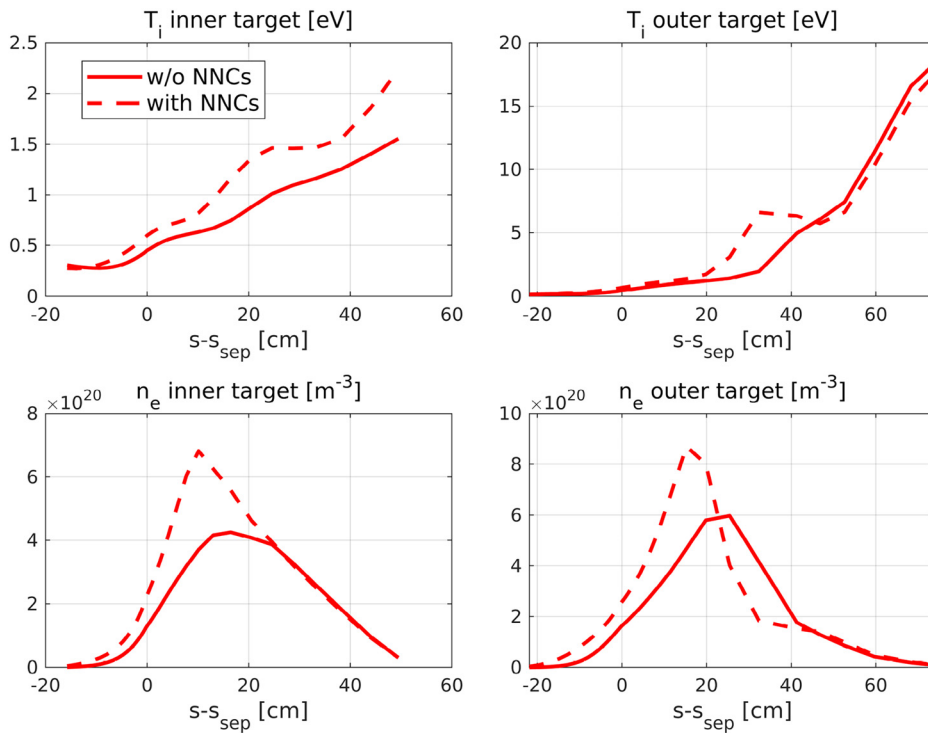


FIG. 4. Electron density and ion temperature along the inner and outer target plate, with and without BGK neutral–neutral collisions, obtained with the most accurate numerical inputs ($P \times 10$ and $\Delta t = 10^{-6}$ s).

deterministic difference depending on P . This is the bias due to the non-linearity of the kinetic equation with NNCs. This bias follows the expected $1/P$ scaling, as shown in Fig. 5(b). However, the bias is small, with only $\sim 10\%$ difference between the case with $P \times 10$ and $P/100$. The differences become even smaller further away from the divertor.

3. Bias error estimation for coupled plasma-neutral simulations

Figure 6 displays the converged and post-processing-averaged target profiles for the case with NNCs with different numerical inputs. Qualitatively, there is convergence to a unique solution for sufficiently high P and sufficiently small Δt . The predicted interdependence of the bias, the time step, and the number of particles is clearly visible in the results. For $\Delta t = 10^{-6}$ s, even the case with $P/100$ is already quite accurate, whereas for $\Delta t = 10^{-4}$ and $\Delta t = 10^{-5}$ s, $P/100$ leads to a large numerical error, pushing the solution to a completely different regime. However, increasing the number of particles drives the solutions with large time steps to the correct solution. Clearly, the same number of particles leads to a lower bias in combination with a smaller time step. The relative differences between the solutions are analyzed quantitatively in Figs. 7 and 8. For $P \rightarrow \infty$, the steady-state result should be independent of Δt . However, here we observe a stagnation toward slightly different solutions. The reason for this is analyzed in Sec. IV A 5.

To summarize the accuracy of each simulation in a concise way, we consider the relative L2-norm error on the electron density over the entire plasma grid

$$\varepsilon_{n_e, L2} = \sqrt{\frac{\sum_i^N V_i \cdot (n_{e,i} - n_{e,REF,i})^2}{\sum_i^N V_i \cdot n_{e,REF,i}^2}}, \quad (10)$$

where N is the number of cells, V_i is the volume of cell i , n_e is the electron density, and $n_{e,REF}$ is the electron density of a reference solution (on the same grid) obtained with a larger number of particles, smaller time step, or both. Because all simulations are performed on the same grid, and because we average over a large number of iterations in the statistical steady state, we can say that $\varepsilon_{n_e, L2}$ is a good approximation for the bias error ε_{bc} .

Figures 7 and 8 summarize the results of the cross-parameter scan on P and Δt . Figures 7(a) and 7(c) shows, for each Δt separately, the relative error compared to the case with $P \times 10$, with and without NNCs, respectively. The bias clearly decreases with $\frac{1}{P}$. There is no noticeable effect of the NNCs on the convergence behavior. When comparing each simulation with the most accurate result ($\Delta t = 10^{-6}$ s and $P \times 10$) in Figs. 7(b) and 7(d), we initially also observe a $\frac{1}{P}$ scaling, but the decrease flattens because of the stagnation toward slightly different solutions. Figures 8(a) and 8(c) shows, for each P separately, the relative error compared to the case with $\Delta t = 10^{-6}$ s, with and without neutral–neutral collisions, respectively. Figures 8(b) and 8(d) again compares each result to the most accurate solution with $\Delta t = 10^{-6}$ s and $P \times 10$. From these figures, it is clear that for a given P the bias increases monotonically with the time step, but the increase is weaker than the theoretical prediction of $\sim \Delta t$ made in Ref. 10. Again, the neutral–neutral collisions do not cause any significant difference in convergence behavior.

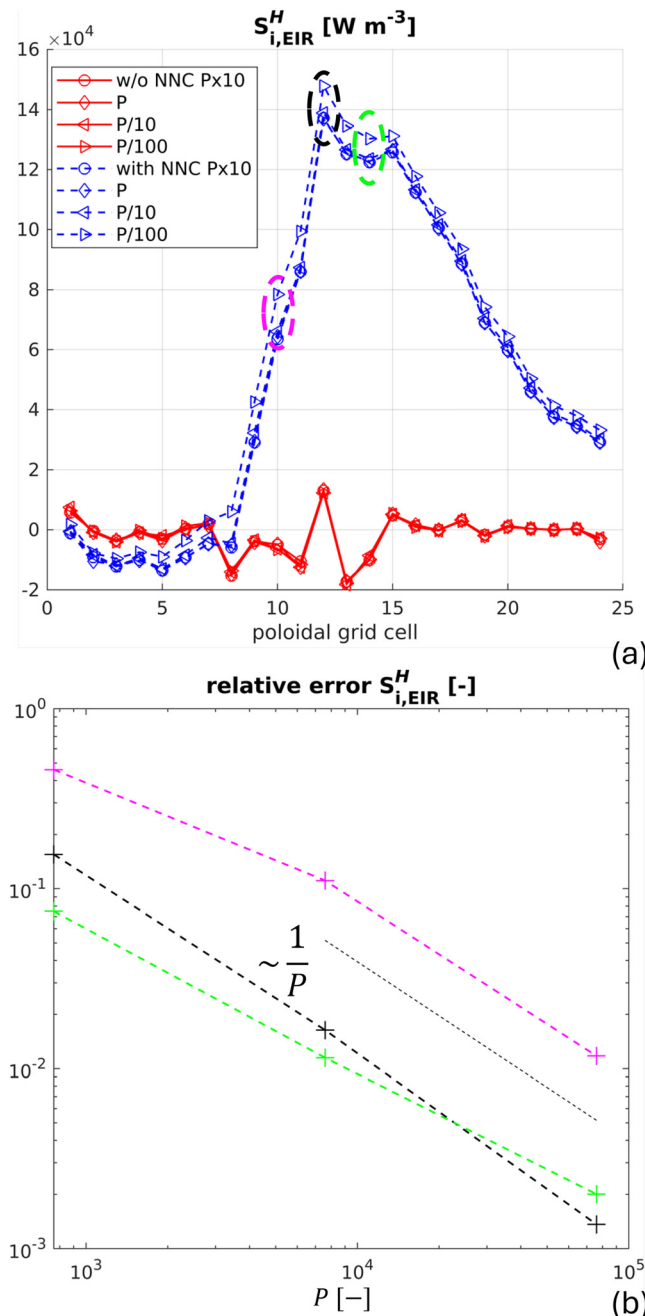


FIG. 5. Ion heat source due to neutral interactions as calculated by EIRENE on a fixed plasma background, in the bottom-most flux tube of the inner divertor leg (a). Relative difference compared to the result with $P \times 10$ for the ion heat source terms encircled in magenta, black, and green (b).

4. Global particle balance

Figure 9 compares the influx and outflux of particles over 1000 iterations in the statistical steady-state phase for the simulations with $\Delta t = 10^{-4}$ s. The influx of particles consists of the gas puff and the

core particle flux. Because the D-only case employs a fixed core density, the resulting core particle flux depends on the solution. In the plot we use the value from the most accurate simulation, but it was checked that the influxes for the other cases differ only a few percent. The figure shows how the global particle balance evolves from being off by an order of magnitude for $P/100$, to being perfectly satisfied for $P \times 10$. It was double-checked that the code guarantees perfect particle conservation for a given plasma background, i.e., for each isonuclear sequence the sum of all EIRENE ionization sources is exactly equal to the sum of the incident fluxes minus the pumped fluxes. The global imbalances arise due to the non-linear plasma-neutral coupling throughout the iterations. For completeness we mention that the target recycling flux is roughly $5.5 \times 10^{24} \text{ s}^{-1}$, a factor 34 larger than the influx.

5. Investigation of convergence anomaly

For the limit of $P \rightarrow \infty$, the steady-state solution should converge toward a unique solution, regardless of the false time step. However, we observe a stagnation in the P -scan toward slightly different solutions for different time steps. Upon deeper investigation, this was traced back to radial oscillations in the parallel ion velocities ($V_{i,\parallel}$). Figure 10 shows a radial cross section of the plasma grid, 8 cells upstream of the inner target (corresponding to ~ 2 cm upstream of the strikepoint along the separatrix), where the encountered anomaly is most pronounced. Although the profiles of n_e (shown in the figure) and temperatures (not shown in the figure) are smooth throughout the iterations, there are radial oscillations in the parallel ion velocity. These oscillations are random in time, because they disappear when averaging the result. Furthermore, the magnitude of the oscillations increases quickly with Δt , but decreases only slowly with P . This is surprising, because it was checked that a deterministic fluid neutral case on the same grid with the same time step does not lead to such oscillations. We appear to have encountered a numerical oscillation that is triggered by any amount of statistical noise, but does not linearly decrease with the magnitude of the statistical noise. This causes the offset in the solution for different time steps for $P \rightarrow \infty$. A possible underlying cause is that the EIRENE sources for parallel momentum are much more noisy than those for particles and energy, because EIRENE uses a collision estimator instead of a track-length estimator for the charge-exchange collisions.

B. D + He artificial slab case

We again perform a cross-parameter scan over Δt and P for the D + He artificial slab case. The effect of the underrelaxation scheme of Sec. II G is studied in detail here. To keep the number of simulations tractable, the underrelaxation is only performed for the cases with $P/100$. Section IV B 1 investigates the impact of the underrelaxation scheme on the convergence speed, Sec. IV B 2 discusses the cross-parameter convergence scan, and the global particle balances are studied in Sec. IV B 3. These first three subsections consider only the case without NNCs, while a subset of the analysis is repeated with NNCs in Sec. IV B 4.

1. Impact of underrelaxation on convergence speed and stability

A first question is whether the underrelaxation scheme has an impact on the convergence time. Figure 11 shows the time evolution

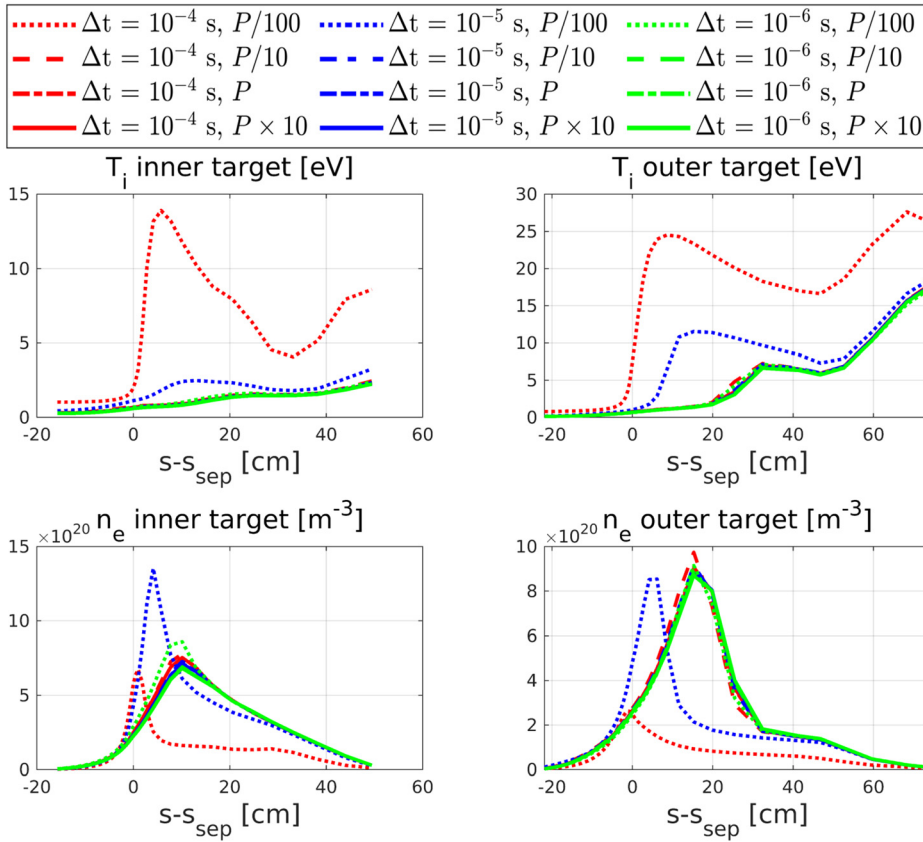


FIG. 6. Ion temperature and electron density along the inner and outer target plate, obtained with the D-only plasma model with NNCs, with different combinations of Δt and P . The cases with smaller time steps and more particles are visually indistinguishable because they converge toward a qualitatively unique solution.

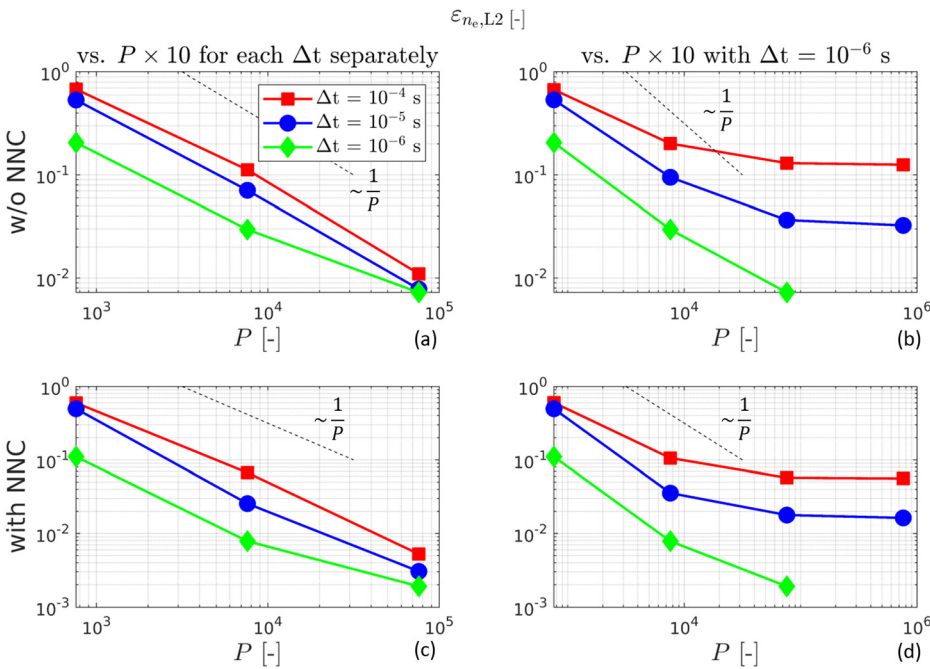


FIG. 7. Relative L2-norm errors of n_e compared to different reference solutions. The top row contains the cases without NNCs and the bottom row those with NNCs included.

09 March 2026 11:53:42

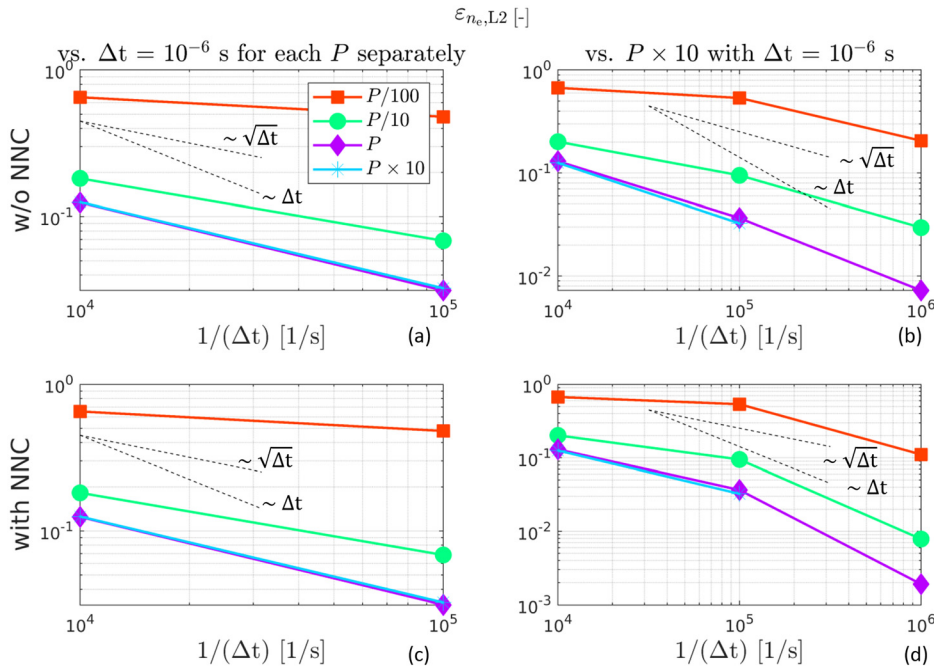


FIG. 8. Relative L2-norm errors of n_e compared to different reference solutions. The top row contains the cases without NNCs and the bottom row those with NNCs included.

starting from a flat initial state ($T_e = T_i = 100$ eV, $n_{D^+} = 10^{19}$ m $^{-3}$, $n_{He^+} = n_{He^{2+}} = 10^{16}$ m $^{-3}$, and zero parallel velocity). Clearly, the underrelaxation parameter N_{ur} does not cause any significant delay of the transient phase.

A second question is to what extent the underrelaxation can be increased. As N_{ur} increases, the Monte Carlo noise reduces, but the time-lag between the underrelaxed sources and the most recent plasma state increases. This time-lag can lead to oscillations and divergence.

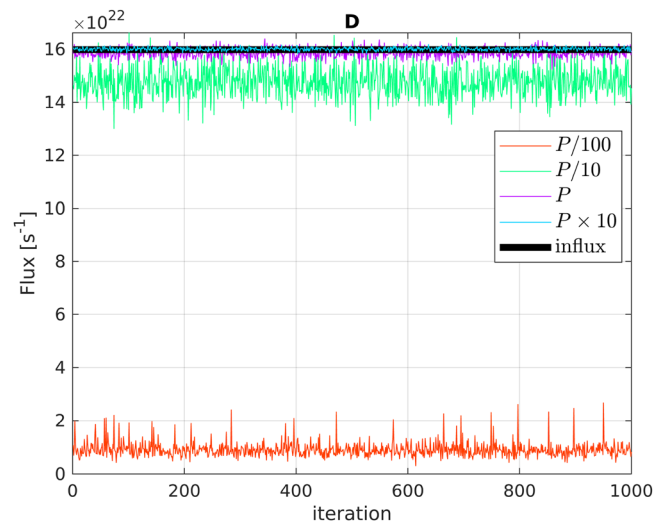


FIG. 9. Comparison of the influx (thick black line) to the outflux of particles (thin colored lines) for different numbers of MC particles, over 1000 iterations in the statistically steady-state phase. Case: $\Delta t = 10^{-4}$ without NNCs.

Table II gives an overview of which cases worked and which cases diverged. These findings indicate that there is roughly an inverse proportionality between the time step and the maximum allowable N_{ur} . This is expected, as the product of Δt and N_{ur} is proportional to the time-lag between the averaged sources and the plasma state.

2. Convergence analysis

Figure 12 displays the ion temperature and electron density along the target for the cross-parameter scan of the cases without underrelaxation. The findings are qualitatively similar to those for the D-only DEMO case, although the magnitudes of the largest bias errors are smaller. For the larger time steps, the cases with $P/100$ end up in a different solution, while for $\Delta t = 10^{-6}$ s, the case with $P/100$ is already quite close to the most accurate result. The anomalous radial oscillations on the parallel velocity as observed for the DEMO D-only case do not occur here, likely due to the simple geometry and high mesh quality of the artificial slab. Therefore, the cross-parameter scan now does converge to a truly unique solution.

Figure 13 provides an example of how the underrelaxation can reduce the noise and the bias. The figure shows a timetrace of over 1000 iterations of the electron temperature at the strikepoint in the steady-state phase of the simulation. For $\Delta t = 10^{-4}$ s and $\Delta t = 10^{-5}$ s, the standard cases with $P/100$ and P are compared to the case with $P/100$ combined with $N_{ur}=100$. The underrelaxation scheme successfully reduces the noise and brings the averaged solution much closer to the correct result, without any increase in CPU time or memory.

We again use the relative L2-norm of the electron density to estimate the bias error. Figure 14(a) gives the estimated bias error based on the case with $P \times 10$, for each Δt separately. Figure 14(b) uses the most accurate result as the reference for all cases. It is again confirmed that the bias decreases with $1/P$.

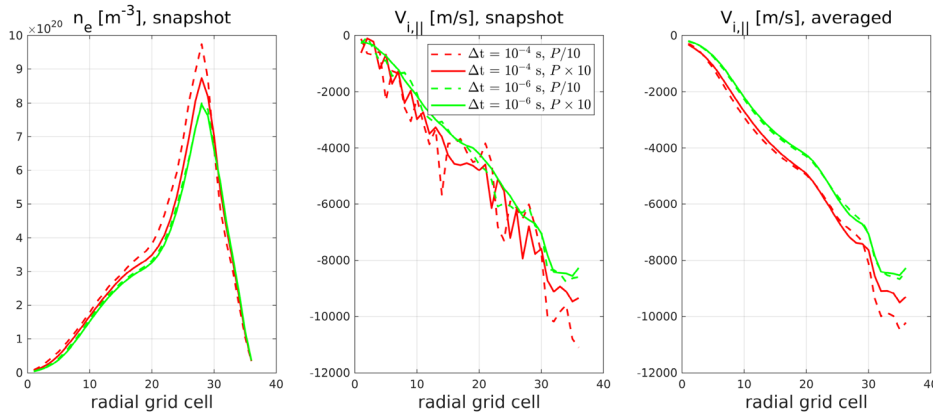


FIG. 10. Radial profiles of n_e and $V_{i,||}$ 8 cells in front of the inner target plate, corresponding to ≈ 2 cm upstream of the inner strikepoint. The left and middle panel shows snapshots of n_e and $V_{i,||}$ in the statistical steady state. The right panel shows the time-averaged $V_{i,||}$ profiles. The profiles are shown for $\Delta t = 10^{-4}$ s (red) and $\Delta t = 10^{-6}$ s (green), for $P/10$ (dashed lines) and $P \times 10$ (full lines).

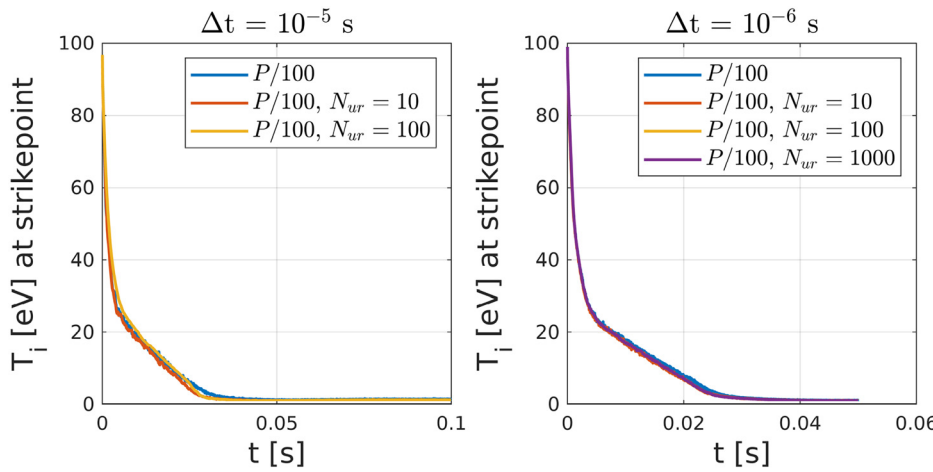


FIG. 11. Time evolution of the ion temperature at the strikepoint, starting from a flat initial state, for different values of N_{ur} .

In Fig. 15, the convergence is plotted as a function of Δt . Similarly as for the D-only DEMO cases, there is for a given P a monotonic decrease in bias for decreasing Δt , but the decrease is generally weaker than $\sim \Delta t$. Figure 15(b) also includes the estimated bias errors for the cases with underrelaxation. The underrelaxation succeeds in reducing the bias by 1–2 orders of magnitude compared to the cases with $P/100$ without underrelaxation [triangles vs red squares in Fig. 15(b)]. However, the cases with underrelaxation exhibit less clear trends with respect to the time step and N_{ur} . A possible explanation is

TABLE II. Convergence results for different combinations of time step and N_{ur} . Green: successful convergence, also when starting from a flat state. Red: divergence, even when starting from an almost-converged state. Orange: success from an almost-converged state, but crash from a flat state. Yellow: successful convergence from a almost-converged state. When starting from a flat state, a temporary increase in the minimum allowable temperatures was needed to prevent a crash.

	$\Delta t = 10^{-4}$ s	$\Delta t = 10^{-5}$ s	$\Delta t = 10^{-6}$ s
$N_{ur} = 10$	Green	Green	Green
$N_{ur} = 100$	Orange	Green	Green
$N_{ur} = 1000$	Red	Red	Yellow

that if N_{ur} gets too close to the stability limit, oscillations of larger magnitude arise, which again reduce the accuracy.

3. Global particle balances

This section studies the convergence of the global particle balances and the effect of the underrelaxation on the balances. Figure 16 compares for a subset of cases the influx of D and He particles to the outfluxes over 1000 iterations in the statistical steady state. Both for D and He, the global balances converge to the correct solution for increasing P . The pumped flux calculated by EIRENE is not explicitly averaged through the underrelaxation scheme, so the statistical variance for $P/100$ is of the same magnitude with and without underrelaxation. However, the averaged result with underrelaxation is much closer to the correct value due to the decreased bias error.

Finally, the relative imbalances for all cases are given in Table III. These results show how the global balances generally improve with increasing number of MC particles, decreasing time step, and increasing underrelaxation. By comparing the global particle imbalances from Table III to the estimated bias errors in Fig. 15, one can see that they are qualitatively related. These findings confirm that tracking the global balances is useful. Although they do not provide information on the precise magnitude of local error contributions, they give a crude




FIG. 12. Ion temperature (left) and electron density (right) along the target plates of the D + He artificial slab case, for different combinations of the number of MC particles P and the time step Δt . The cases with smaller time steps and more particles become visually indistinguishable because they converge toward a unique solution.

indication about the numerical convergence, without needing a computationally more expensive reference solution.

4. Inclusion of neutral-neutral collisions

When including BGK neutral-neutral collisions in the slab case, the target ion temperature changes about 5%. This is a smaller effect compared to the D-only case (Fig. 4), but still sufficient to determine if the NNCs have an effect on the convergence behavior. We now

include a single case with underrelaxation ($\Delta t = 10^{-5}$ s, $P/100$, and $N_{ur} = 100$). Notice that for now, the underrelaxation scheme is only applied to the plasma source terms and not yet to the BGK neutral background. The latter can be considered for completeness, but the results of this paper indicate that this is not a priority. Figure 17 displays the target ion temperatures for a subset of the combinations of Δt and P and for the case with underrelaxation. Again, the interdependence of the time step and bias is clearly visible. The underrelaxation scheme (magenta dashed line) succeeds in reducing the bias, leading to




FIG. 13. Timetraces of the electron temperature at the strikepoint over 1000 iterations in the statistical steady state for different cases.

09 March 2026 11:53:42

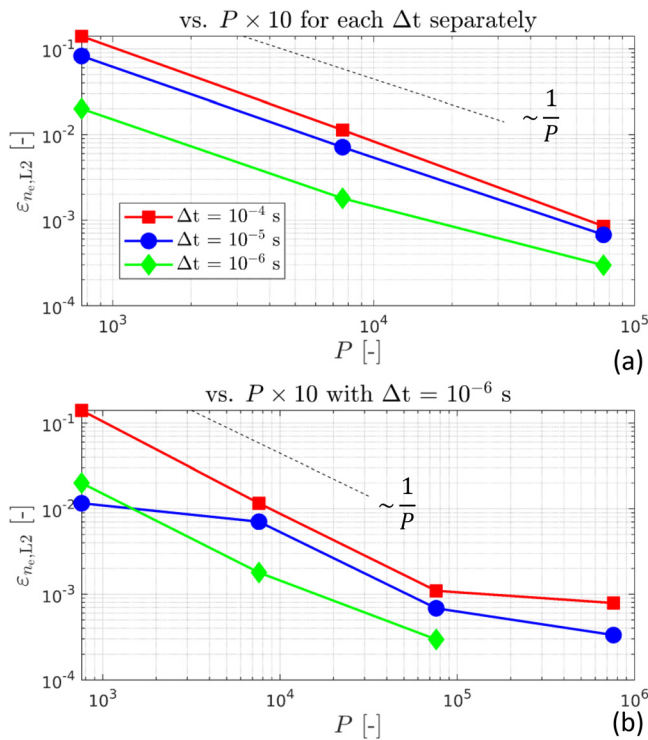


FIG. 14. Relative L2-norm errors of n_e compared to different reference solutions.

a solution that is almost identical to the numerically most accurate ones, while simulating only very few particles per iteration. This is further clarified in Fig. 18, which shows the relative L-2 norm error on n_e . Comparing Figs. 18 and 14(a), there is no noticeable impact of NNCs on the convergence behavior. Because the bias error due to NNCs is negligible compared to that of the plasma–neutral coupling, the under-relaxation scheme achieves a large bias reduction, although the BGK neutral background itself is not averaged.

C. D + He + Ar DEMO case

Because of the high computational cost of the multi-species cases, we study only the runs with $\Delta t = 10^{-5}$ s for $P/10$, P , and $P \times 10$, one case with underrelaxation: $\Delta t = 10^{-5}$ s, $P/10$, $N_{ur} = 100$, and the case with $\Delta t = 10^{-6}$ s and $P/10$. Because of the limited number of cases it is not possible to present a quantitative convergence assessment, but the findings of Secs. IV A and IV B can be qualitatively confirmed.

Initially, the number of particles per stratum for He and Ar were taken the same as for D. However, it was then observed that the bias errors were up to two orders of magnitude larger than for the D-only case. Because the average EIRENE CPU time per particle is a factor 4 lower for He than for D, and a factor 12 lower for Ar than for D for these cases, it was decided to increase the number of impurity particles with a factor 10, as shown in Table I. This partially resolved the issue. The case indicated with “IMP/10” in Fig. 19 provides an example of the very large numerical errors that were observed when the number of impurity particles were not multiplied with a factor 10 compared to deuterium.

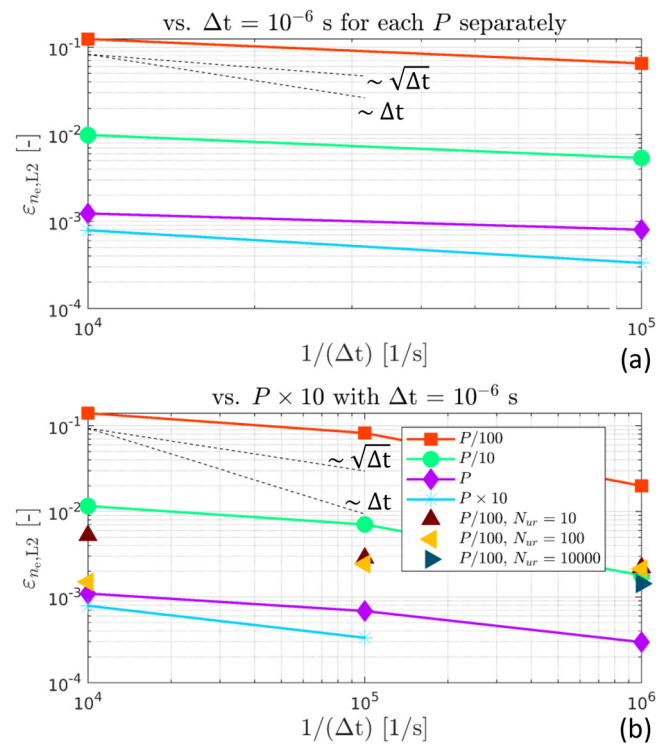


FIG. 15. Relative L2-norm errors of n_e compared to different reference solutions.

1. Target profiles

Figure 19 shows the target profiles corresponding to the different numerical settings. The inner target is deeply detached, with ion temperatures below 1.6 eV along the full length of the target plate. The outer target is partially detached, with an ion temperature around 1 eV at the strikepoint but increasing to above 20 eV toward the baffle. Several conclusions can be made about the numerical convergence. Since the cases with P and $P \times 10$ for $\Delta t = 10^{-5}$ differ less than 5%, we can safely assume that $P \times 10$ is close to convergence. The interdependence of the time step and bias is again proven, as the case with $P/10$ is closer to the correct solution with $\Delta t = 10^{-6}$ s than with $\Delta t = 10^{-5}$ s. Comparing Fig. 19 with Fig. 6, the bias error for $P/10$ and $\Delta t = 10^{-5}$ s is much larger for the high-power multi-species cases than for the low-power D-only case, although we had already increased the number of impurity MC particles by a factor of 10. On the outer target, the underrelaxation scheme successfully reduces the bias. On the inner target, however, there remains an offset. It is expected that here the underrelaxation scheme is close to the stability limit.

2. Particle balances

Figure 20 shows how increasing the number of MC particles per iterations generally improves the global particle balance. Similarly, as before, the underrelaxation scheme improves the balance as well. However, contrary to the D-only DEMO case (Fig. 9) and the slab case (Fig. 16), the outfluxes for He and Ar now seem to converge toward a slightly wrong value, corresponding to an imbalance of 5%–6%. This was

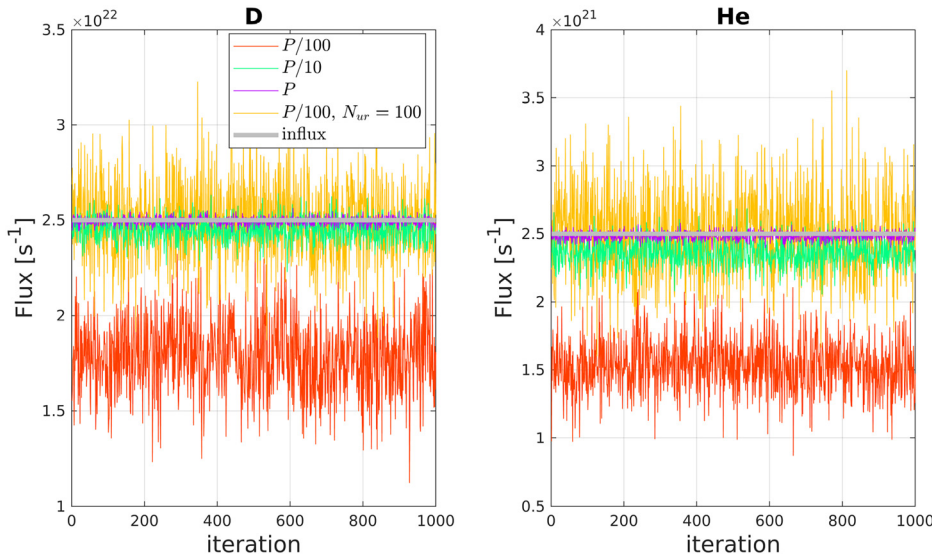


FIG. 16. Comparison of particle influxes (thick gray lines) to the particle outfluxes (thin colored lines) for deuterium (left) and helium (right), over 1000 iterations in the statistical steady-state phase for $\Delta t = 10^{-4}$ s.

traced back to limits imposed on the density and parallel velocity. In SOLPS-ITER, it is a common problem that the parallel velocities of certain impurities achieve unrealistically high values in regions where their density is low. These exuberant parallel velocities can further lead to artificial viscous heating, eventually leading to the crashing of the code. For now, this is resolved by imposing a minimum density n_{\min} and maximum Mach number M_{\max} . In the set-up of the multi-species DEMO cases, $n_{\min} = 10^4 \text{ m}^{-3}$ and $M_{\max} = 2$ were selected. It was found that for

every ionization stage of argon, the minimum particle density limit is reached in part of the domain. Additionally, it was found that the Mach 2 limit is reached in a substantial part of the divertor legs for the helium ions. A subset of these findings is displayed in Fig. 21. When such user-imposed limits are reached, the global balances cannot be satisfied perfectly, even for $P \rightarrow \infty$. Mainly, the choice of $M_{\max} = 2$ was, in hindsight, too restrictive.

TABLE III. Relative global particle imbalances for the D + He slab cases. The precise values below 0.1% are not listed because they are dominated by residual statistical noise.

	D global particle imbalance (%)		
	$\Delta t = 10^{-4}$ s	$\Delta t = 10^{-5}$ s	$\Delta t = 10^{-6}$ s
$P/100$	28.6	15.4	1.6
$P/10$	2.1	0.8	0.25
P	0.2	0.13	<0.1
$P \times 10$	<0.1	<0.1	<0.1
$P/100, N_{ur} = 10$	0.26	0.22	<0.1
$P/100, N_{ur} = 100$	0.13	0.12	<0.1
$P/100, N_{ur} = 1000$	<0.1

	He global particle imbalance (%)		
	$\Delta t = 10^{-4}$ s	$\Delta t = 10^{-5}$ s	$\Delta t = 10^{-6}$ s
$P/100$	38.5	26.2	11.4
$P/10$	5.6	3.08	1.02
P	0.65	0.41	0.29
$P \times 10$	0.15	0.11	0.16
$P/100, N_{ur} = 10$	0.87	0.82	1.0
$P/100, N_{ur} = 100$	<0.1	0.19	0.5
$P/100, N_{ur} = 1000$	<0.1

D. Accuracy vs computational cost

A detailed optimization of total CPU cost vs total numerical accuracy is not performed here. A methodology to do this was already




FIG. 17. Ion temperature along the target plates of the D + He artificial slab case, for different combinations of the number of MC particles P and the time step Δt , and one case with underrelaxation. Case with NNCs.




FIG. 18. Relative L2-norm errors of n_e for the D + He slab case with NNCs.

proposed in Ref. 8. In this paper, we limit ourselves to some useful pointers.

The increase in bias with Δt does not mean that it is generally advised to reduce the time step, because the required number of

iterations scales with $\frac{1}{\Delta t}$, linearly increasing the total CPU cost required by both EIRENE and B2.5. In contrast, increasing P only effects the EIRENE CPU time. Furthermore, the EIRENE wall time can be easily reduced by parallelizing over a large number of cores. However, the time step is often limited due to the stability of the plasma model, particularly when drifts and currents are included. Without drifts, larger time steps are typically possible, but a too large time step can still result in numerical issues, such as the oscillations on the parallel velocity that were observed here. An important insight is that, when one is forced to reduce Δt due to the stability of the plasma model, one can reduce P , while still retaining the desired accuracy. The effect of the non-linearity of the neutral-neutral collisions on the convergence behavior was found to be negligible compared to the non-linearity of the plasma-neutral coupling, hence this effect must not be taken into account when optimizing the CPU cost. The underrelaxation scheme introduces an additional dimension to the optimization problem. The stability limit of the underrelaxation, which in itself depends on the time step, further complicates matters. Based on the results of this paper we recommend for now $N_{ur} = 10$ for $\Delta t = 10^{-4}$ s and $N_{ur} = 100$ for $\Delta t = 10^{-5}$ s or $\Delta t = 10^{-6}$ s. The stability limit is likely case dependent, so more research is needed to find generalized recommendations for N_{ur} .

V. CONCLUSIONS AND FUTURE WORK

The numerical convergence behavior of SOLPS-ITER plasma boundary simulations with kinetic neutrals was studied by performing a cross-parameter scan on the time step Δt and the number of Monte

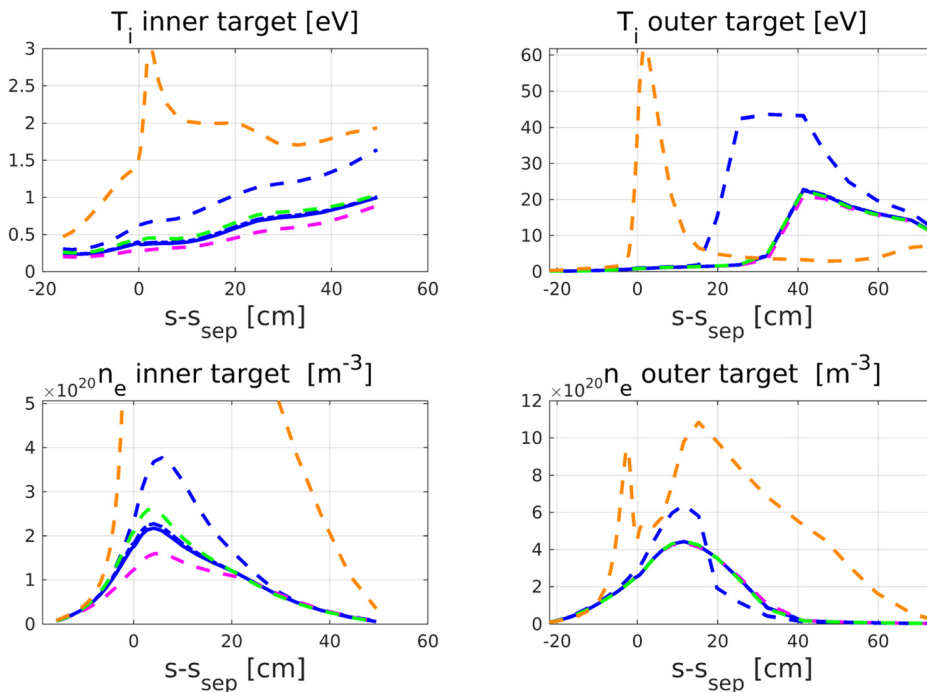



FIG. 19. Ion temperature and electron density along the inner and outer target plate, obtained with the D + He + Ar plasma model with NNCs for different numerical parameters.

09 March 2026 11:53:42

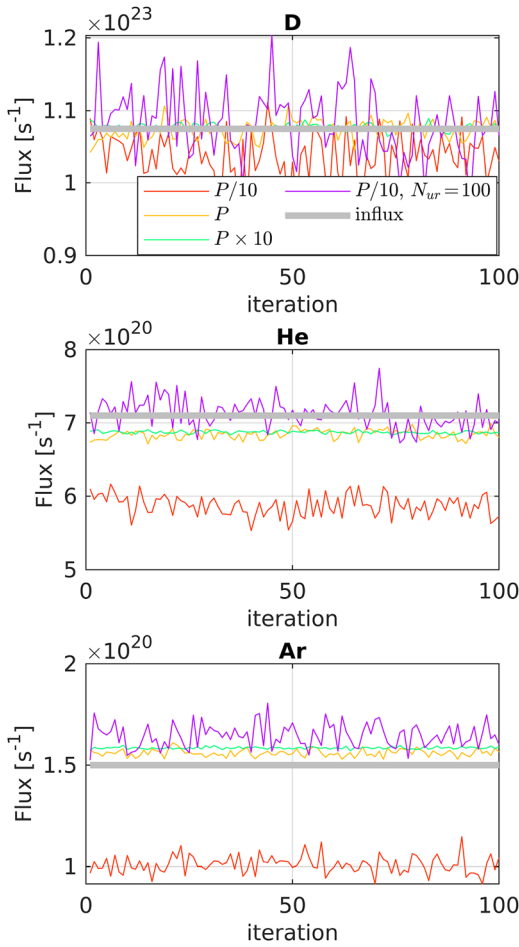


FIG. 20. Comparison of particle influxes (thick gray lines) to the particle outfluxes (thin colored lines) for deuterium (top), helium (middle), and argon (bottom), for $\Delta t = 10^{-5}$ s. The fluxes are shown over 100 iterations in the statistical steady-state phase.

Carlo particles per iteration P . This was done for a low-power deuterium-only DEMO case, a D + He artificial slab, and a D + He + Ar high-power DEMO case. The paper focused on the bias error, which is the deterministic offset in the expected value of the solution upon convergence to a statistical steady-state. The interdependence of Δt , P , and the bias error was clearly exposed. For a given time step, the bias error reduces with $1/P$, as expected from theory. For a given P , the bias error increases monotonically with Δt , but the precise scaling is less clear, typically varying between Δt and $\sqrt{\Delta t}$. This knowledge can be exploited when one is forced to take a small time step on account of the plasma model, for example when drifts and currents are included. Inversely, the combination of a large time step and a low number of particles produces large numerical errors. Increasing the time step from $\Delta t = 10^{-7}$ to $\Delta t = 10^{-5}$ s (for example, when switching from a case with drifts to a case without drifts) will increase the bias error by a factor 10–100 if the number of particles per iteration is kept the same.

It was shown that BGK neutral–neutral collisions lead to an additional component to the bias error. However, this bias contribution

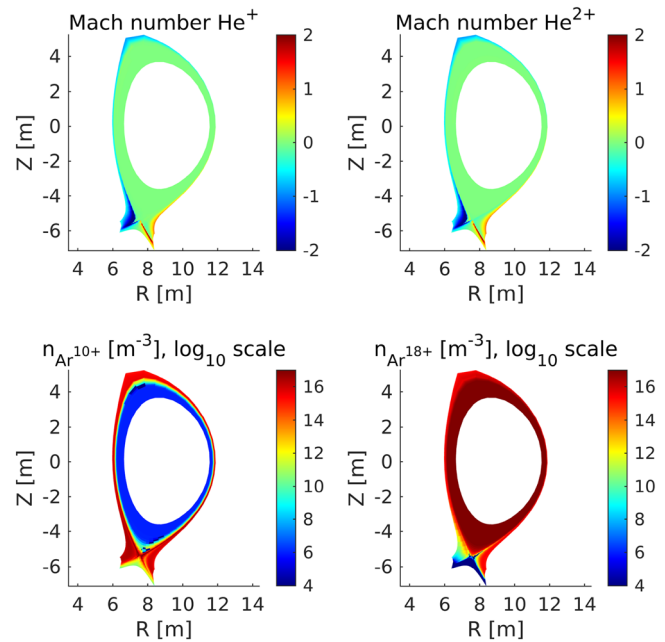


FIG. 21. Mach numbers of the parallel velocities (including the sign) of the He^+ (top left) and He^{2+} (top right) ions, and particle densities of the Ar^{10+} (bottom left) and Ar^{18+} (bottom right) ions. The displayed results are for the case with $\Delta t = 10^{-5}$ s and $P \times 10$.

appeared to be negligible compared to the one arising from the plasma–neutral coupling for the studied cases. If in the future cases are discovered where the bias contribution due to neutral–neutral collisions is not negligible, then the underrelaxation scheme should be implemented for the neutral BGK background as well.

On an identical DEMO geometry, the bias errors for the high-power D + He + Ar case were found to be roughly one order of magnitude larger than for a low-power D-only case, although the number of MC particles for the impurities themselves was already increased by a factor 10. The reason is not precisely known; possible explanations are that the Monte Carlo techniques used by EIRENE are not optimized for absorption-dominated argon neutrals, or that stronger non-linearities are present in the plasma–neutral coupling in the more reactor-relevant case. The large bias errors encountered for the D + He + Ar DEMO case prove that the study of this numerical error remains highly relevant.

The performance of an underrelaxation scheme for the noisy kinetic neutral source terms was investigated. The scheme appears promising to speed up the convergence of SOLPS-ITER simulations, at least for steady-state cases. It was shown that the number of iterations needed to reach steady state does not increase when using the scheme with a fixed underrelaxation parameter (contrary to the Robbins–Monro approach, where the neutral sources are averaged over all iterations). The scheme successfully reduced the bias errors by one to two orders of magnitude, without requiring any additional computational cost or data storage. However, it was shown that the underrelaxation scheme can become unstable or reduce accuracy if the product of the time step and the underrelaxation parameter is too high. A deeper understanding of the stability of the scheme should be the goal of

future work. The performance of the scheme should be tested for more cases. An important shortcoming of the presented work is the absence of drifts. Drifts and currents are not expected to affect the scaling of the numerical errors, but they can affect the absolute values. In principle, the underrelaxation scheme can be applied to drift cases without change, but the stability could be affected. Investigation of drift cases should be a priority for future research. The underrelaxation scheme should then also be compared to the use of inner iterations, which is currently the most popular strategy to reduce the computational cost of EIRENE for drift cases.^{23,24}

For now, we have used a spatially constant time step (except for the core multipliers). When searching for a steady-state solution, this is not required.²³ The effect of spatially varying timesteps on the bias errors and on the stability of the underrelaxation scheme should also be investigated in future work.

It was shown that the global particle balances are qualitatively linked to the bias errors and can hence be used to get a rough idea about the numerical convergence of a simulation case. For the D-only DEMO case and for the D+He slab case, global imbalances of less than 0.1% were achieved with the numerically most accurate cases. The underrelaxation scheme improves the balance. For the D+He+Ar DEMO case, global particle imbalances of 5%–6% remained for the most accurate case. This was traced back to user-imposed limits on the minimum particle density and maximum Mach number.

Finally, it should be investigated why the observed scaling of the bias error with Δt in the presented SOLPS-ITER simulations does not precisely follow the theoretical predictions made in previous theoretical work.¹⁰ This can be done by starting from the 1D simulations of Ref. 10 and step-by-step elaborating the physical and geometric description toward that of SOLPS-ITER.

ACKNOWLEDGMENTS

Parts of this work have been carried out within the framework of the EUROfusion Consortium, funded by the European Union via the Euratom Research and Training Programme (Grant Agreement No. 101052200—EUROfusion). Views and opinions expressed are, however, those of the author(s) only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the European Commission can be held responsible for them. N. Horsten is a postdoctoral research fellow of the Research Foundation Flanders (FWO) under Grant No. 12AES24N. 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.

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

Author Contributions

Wim Van Uytven: Conceptualization (lead); Formal analysis (lead); Methodology (lead); Software (lead); Writing – original draft (lead). **F. Subba:** Data curation (equal); Funding acquisition (supporting);

Supervision (equal); Writing – review & editing (equal). **S. Wiesen:** Data curation (equal); Funding acquisition (equal); Software (equal); Supervision (equal); Writing – review & editing (supporting). **N. Horsten:** Conceptualization (equal); Software (equal); Supervision (supporting); Writing – review & editing (equal). **Z. Tang:** Investigation (equal); Methodology (equal); Writing – review & editing (equal). **W. Dekeyser:** Conceptualization (equal); Formal analysis (equal); Software (equal); Supervision (lead); Writing – review & editing (equal).

DATA AVAILABILITY

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

REFERENCES

- S. Wiesen, D. Reiter, V. Kotov, M. Baelmans, W. Dekeyser, A. Kukushkin, S. Lisgo, R. Pitts, V. Rozhansky, G. Saibene *et al.*, “The new SOLPS-ITER code package,” *J. Nucl. Mater.* **463**, 480–484 (2015).
- X. Bonnin, W. Dekeyser, R. Pitts, D. Coster, S. Voskoboinikov, and S. Wiesen, “Presentation of the new SOLPS-ITER code package for tokamak plasma edge modelling,” *Plasma Fusion Res.* **11**, 1403102 (2016).
- R. Schneider, D. Coster, B. Braams, P. Xantopoulos, V. Rozhansky, S. Voskoboinikov, L. Kovaltsova, and H. Bürbaumer, “B2-solps5. 0: SOL transport code with drifts and currents,” *Contrib. Plasma Phys.* **40**, 328–333 (2000).
- D. Reiter, M. Baelmans, and P. Börner, “The EIRENE and B2-EIRENE codes,” *Fusion Sci. Technol.* **47**, 172–186 (2005).
- K. Ghoo, W. Dekeyser, G. Samaey, P. Börner, and M. Baelmans, “Accuracy and convergence of coupled finite-volume/Monte Carlo codes for plasma edge simulations of nuclear fusion reactors,” *J. Comput. Phys.* **322**, 162–182 (2016).
- K. Ghoo, P. Börner, W. Dekeyser, A. Kukushkin, and M. Baelmans, “Grid resolution study for B2-EIRENE simulation of partially detached ITER divertor plasma,” *Nucl. Fusion* **59**, 026001 (2019).
- K. Ghoo, G. Samaey, and M. Baelmans, “Accuracy and convergence of iteratively solved Monte Carlo codes for simulations in the plasma edge of nuclear fusion reactors,” *Contrib. Plasma Phys.* **58**, 652–658 (2018).
- K. Ghoo, “Accuracy-based simulation strategies for plasma edge simulations for nuclear fusion devices,” Ph.D. thesis (KU Leuven, 2019).
- W. Van Uytven, W. Dekeyser, F. Subba, S. Wiesen, N. Horsten, N. Vervloesem, and M. Baelmans, “Discretization error estimation for EU-DEMO plasma-edge simulations using SOLPS-ITER with fluid neutrals,” *Contrib. Plasma Phys.* **64**, e202300125 (2024).
- M. Baeten, K. Ghoo, M. Baelmans, and G. Samaey, “Analytical study of statistical error in coupled finite-volume/Monte Carlo simulations of the plasma edge,” *Contrib. Plasma Phys.* **58**, 659–665 (2018).
- P. L. Bhatnagar, E. P. Gross, and M. Krook, “A model for collision processes in gases. I. Small amplitude processes in charged and neutral one-component systems,” *Phys. Rev.* **94**, 511 (1954).
- V. Kotov, D. Reiter, and A. S. Kukushkin, “Numerical study of the ITER divertor plasma with the B2-EIRENE code package,” Report No. 4257 [Forschungszentrum Juelich (Germany). Institute fuer Energieforschung (IEF), 2007].
- V. Kotov and D. Reiter, “B2-B2. 5 code benchmarking, Part III: Convergence issues of the B2-EIRENE code,” Technical Report No. JUL-4371 (Forschungszentrum Juelich GmbH, 2014).
- B. Mortier, M. Baelmans, and G. Samaey, “A comparison of source term estimators in coupled finite-volume/Monte-Carlo methods with applications to plasma edge simulations in nuclear fusion,” *arXiv:2012.08981* (2020).
- H. Kawashima, K. Shimizu, T. Takizuka, S. Sakurai, T. Nakano, N. Asakura, and T. Ozeki, “Development of integrated SOL/divertor code and simulation study in JAEA,” *Plasma Fusion Res.* **1**, 031 (2006).
- H. Robbins and S. Monro, “A stochastic approximation method,” *Ann. Math. Stat.* **22**, 400–407 (1951).
- V. Kotov, “Particle conservation in numerical models of the tokamak plasma edge,” *Phys. Plasmas* **24**, 042511 (2017).

- ¹⁸Y. Marandet, H. Bufferand, G. Ciraolo, P. Genesio, P. Meliga, J. Rosato, E. Serre, and P. Tamain, "Effect of statistical noise on simulation results with a plasma fluid code coupled to a Monte Carlo kinetic neutral code," *Contrib. Plasma Phys.* **56**, 604–609 (2016).
- ¹⁹F. Subba, D. Coster, M. Moscheni, and M. Siccino, "SOLPS-ITER modeling of divertor scenarios for EU-DEMO," *Nucl. Fusion* **61**, 106013 (2021).
- ²⁰M. Blommaert, W. Dekeyser, N. Horsten, P. Börner, and M. Baelmans, "Implementation of a consistent fluid-neutral model in SOLPS-ITER and benchmark with EIRENE," *Contrib. Plasma Phys.* **58**, 718–724 (2018).
- ²¹W. Van Uytven, M. Blommaert, W. Dekeyser, N. Horsten, and M. Baelmans, "Implementation of a separate fluid-neutral energy equation in SOLPS-ITER and its impact on the validity range of advanced fluid-neutral models," *Contrib. Plasma Phys.* **60**, e201900147 (2020).
- ²²W. Dekeyser, P. Börner, S. Voskoboynikov, V. Rozhansky, I. Senichenkov, E. Kaveeva, I. Veselova, E. Vekshina, X. Bonnin, R. Pitts *et al.*, "Plasma edge simulations including realistic wall geometry with SOLPS-ITER," *Nucl. Mater. Energy* **27**, 100999 (2021).
- ²³E. Kaveeva, V. Rozhansky, I. Senichenkov, I. Veselova, S. Voskoboynikov, E. Sytova, X. Bonnin, and D. Coster, "Speed-up of SOLPS-ITER code for tokamak edge modeling," *Nucl. Fusion* **58**, 126018 (2018).
- ²⁴M. Carpita, E. Tonello, C. Colandrea, O. Février, H. Reimerdes, G.-Y. Sun, and F. Mombelli, "Verification of techniques for accelerated and stable SOLPS-ITER simulations including plasma drifts," *Nucl. Fusion* **65**, 116001 (2025).