

















Fluid, kinetic and hybrid approaches for neutral and trace ion edge transport modelling in fusion devices

D.V. Borodin^{1,*} , F. Schluck¹ , S. Wiesen¹ , D. Harting¹, P. Börner¹,
S. Brezinsek¹ , W. Dekeyser² , S. Carli² , M. Blommaert² ,
W. Van Uytven² , M. Baelmans² , B. Mortier³ , G. Samaey³ ,
Y. Marandet⁴ , P. Genesio⁴, H. Bufferand⁵, E. Westerhof⁶ ,
J. Gonzalez⁶ , M. Groth⁷ , A. Holm⁷ , N. Horsten⁷ and H.J. Leggate⁸

¹ Forschungszentrum Jülich GmbH, Institut für Energie- und Klimaforschung – Plasmaphysik, Partner of the Trilateral Euregio Cluster (TEC), 52425 Jülich, Germany

² KU Leuven, Department of Mechanical Engineering, Celestijnenlaan 300, 3001 Leuven, Belgium

³ KU Leuven, Department of Computer Science, Celestijnenlaan 200A, 3001 Leuven, Belgium

⁴ Aix-Marseille Univ., CNRS, PIIM, F-13013 Marseille, France

⁵ CEA, IRFM, F-13108 Saint-Paul-lez-Durance, France

⁶ DIFFER—Dutch Institute for Fundamental Energy Research, Partner in the Trilateral Euregio Cluster, De Zaale 20, 5612 AJ Eindhoven, Netherlands

⁷ Department of Applied Physics, Aalto University, Espoo, Finland

⁸ National Centre for Plasma Science and Technology, Dublin City University, Glasnevin Dublin 9, Ireland

E-mail: d.borodin@fz-juelich.de

Received 16 July 2021, revised 24 November 2021

Accepted for publication 3 December 2021

Published 7 July 2022




CrossMark

Abstract

Neutral gas physics and neutral interactions with the plasma are key aspects of edge plasma and divertor physics in a fusion reactor including the detachment phenomenon often seen as key to dealing with the power exhaust challenges. A full physics description of the neutral gas dynamics requires a 6D kinetic approach, potentially time dependent, where the details of the wall geometry play a substantial role, to the extent that, e.g., the subdivertor region has to be included. The Monte Carlo (MC) approach used for about 30 years in EIRENE (Reiter *et al* 2005 *Fusion Sci. Technol.* **47** 172–86), is well suited to solve these types of complex problems. Indeed, the MC approach allows simulating the 6D kinetic equation without having to store the velocity distribution on a 6D grid, at the cost of introducing statistical noise. MC also provides very good flexibility in terms of geometry and atomic and molecular (A & M) processes. However, it becomes computationally extremely demanding in high-collisional regions (HCRs) as anticipated in ITER and DEMO. Parallelization on particles helps reducing the simulation wall clock time, but to provide speed-up in situations where single trajectories potentially involve a very large number of A & M events, it is important to derive a hierarchy of models in terms of accuracy and to clearly identify for what type of physics issues they provide reliable answers. It was demonstrated that advanced fluid neutral models are very accurate in HCRs, and at least an order of magnitude faster than fully kinetic simulations. Based on these fluid models, three hybrid fluid–kinetic approaches are introduced: a spatially

* Author to whom any correspondence should be addressed.

 Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

hybrid technique, a micro–macro hybrid method, and an asymptotic-preserving MC scheme, to combine the efficiency of a fluid model with the accuracy of a kinetic description. In addition, A & M ions involved in the edge plasma chemistry can also be treated kinetically within the MC solver, opening the way for further hybridisation by enabling kinetic impurity ion transport calculations. This paper aims to give an overview of methods mentioned and suggests the most prospective combinations to be developed.

Keywords: EIRENE code, neutrals, Monte Carlo, code performance, transport simulations, fluid–kinetic hybridisation, edge and divertor plasma

(Some figures may appear in colour only in the online journal)

Introduction

Numerical simulations with a kinetic code such as EIRENE [1] are indispensable for both understanding and predicting the fuel and impurity transport in the edge and divertor areas of fusion devices. The transport determines impurity penetration towards the core, plasma exhaust and plasma–surface interaction (PSI) issues having impact on the duty cycle of ITER and DEMO. The insight into the interplay of transport and atomic–molecular (A & M) processes [2, 3] provided by modelling is key for understanding of the detachment phenomenon [4, 5], which is critical for many exhaust regimes envisaged for ITER and DEMO.

EIRENE is a multi-purpose Boltzmann-equation Monte Carlo (MC) solver typically employed in an iterative scheme with a fluid plasma code, succinctly referred to as computational fluid dynamics (CFD) code in this paper, see figure 1. A number of 2D CFD-EIRENE code packages such as SOLPS-ITER [6, 7], EDGE2D-EIRENE [8], and SOLEdge2D-EIRENE [9] are extensively employed by the fusion community. Also 3D codes are applied, such as EMC3-EIRENE [10] or SolEdge3X-EIRENE [11], and are more demanding in terms of CPU and memory requirements. All the mentioned kinetic–fluid packages are under active development. They provide self-consistently generated plasma distributions of macroscopic properties (2D or 3D), heat and particle fluxes to the wall, synthetic spectroscopy and radiative losses. An essential part of the neutral MC tracing procedure are the collisional–radiative models (CRMs) for A & M processes involving main-plasma species and intrinsic/extrinsic impurities. This includes ionization–dissociation–recombination of A & M species, break-up chains of molecules in plasma, and elastic processes.

The plasma fluid side of the packages is typically much less CPU time demanding than the EIRENE side. However, due to the strong non-linear coupling between neutrals and plasma, a large number of iterations of the coupled code systems is needed to reach a statistically stationary steady state. Despite hundreds or even thousands of CFD-EIRENE iterations are typically necessary to converge to a steady state solution, all the established packages mentioned above mostly provide sufficient performance allowing transport calculations for current devices. The situation becomes more difficult for ITER [12],

DEMO and other large devices, for which the equilibration of the balance equations becomes much slower. In particular the particle balance poses a bottleneck, because the time scale for the particle balance τ_p^* scales as $1/(1 - R)$ [13], with the effective particle recycling coefficient R approaching unity for ITER and DEMO configurations. The time constant τ_p^* may be in the range of seconds for ITER and DEMO, while time steps in the simulations are typically restricted to the order of 10^{-6} s or lower due to numerical stability constraints including stiffness on the CFD side of the package. Energy and momentum have much lower effective recycling coefficients, leading to much faster equilibration and posing less of a constraint. Also important is the level of approximation, for instance turbulent transport is often modelled using the gradient diffusion approximation, because full turbulence simulations are much more computationally demanding [11]. In fact, even certain experiments at relatively small linear plasma devices can be challenging for EIRENE simulations for instance if time-dependent simulations are necessary, whenever the time step imposed by the plasma solver is small compared to the lifetime of the particles as is typical for highly dissipative detached conditions [12]. A good example of that are transient studies of the lithium vapour box experiment at Magnum-PSI aimed at detachment [14].

Generally, plasma ions and electrons are simulated by a fluid approach, and only the neutrals including molecular species are treated kinetically or in hybrid approaches discussed in section 1. In fusion-relevant plasmas, in major parts of the domain the mean free path λ for neutrals is large compared to the gradient lengths L and the neutral flow is in the large Knudsen number regime, $K_n = \lambda/L$, for which no accurate fluid closure is available. Hence, a kinetic approach is generally used for neutrals. However, especially in large machines such as ITER and in detached regimes, high-collisional regions (HCRs), where K_n becomes low, may appear. In these HCRs, the coupling of the neutrals with the background plasma becomes very strong, leading to quasi-Maxwellian distributions for neutrals through e.g. charge exchange reactions with background ions. Typical charge exchange mean-free paths for ITER can be in the mm-range in HCRs in the divertor to m-range in the main chamber, varying over several orders of magnitude mainly due to the inverse dependence on the plasma density. Also neutral–neutral collisions can become

recombination, gas puff and also sources of atoms stemming, e.g., from dissociation of molecules are included in this term. The second term on the right-hand side, $S_v(\mathbf{x}, \mathbf{v}, t)$ ($f(\mathbf{x}, \mathbf{v}, t)$), describes the impact of collision/absorption events such as charge-exchange (CX) scattering, ionization, and atom reflection at boundaries, which all depend linearly on the distribution function $f(\mathbf{x}, \mathbf{v}, t)$. For brevity of notation, we will omit the arguments $\mathbf{x}, \mathbf{v}, t$ in what follows. Using the linearity of the equation, we can write $f = \sum_i f_i$, where f_i is the solution for source $S_{c,i}$. Nonlinear neutral–neutral collisions following a BGK [19] approach can be included iteratively in this framework. This is subject of ongoing research, however outside the focus of the present paper. The basic fluid–kinetic hybridization (FKH) methodology will remain the same when including neutral–neutral collisions in a BGK framework, but approaches relying on the linearity of the problem need to be generalized. It should also be noted that neutral–neutral collisions provide an extra relaxation mechanism towards local equilibrium for neutrals, and would thus tend to reduce the Knudsen number (that is, make the neutrals ‘more fluid’). However, while these collisions are important for the overall model accuracy, they are not necessarily dominant in the same spatial regions where CX-collisions lead to the computational bottleneck. This is subject of ongoing research.

1.1. Advanced fluid neutral models (AFN)

A fluid approach is obtained by taking moments of the kinetic equation:

$$\int_{\mathbf{v}} \mu(\mathbf{v}) \left(\frac{\partial f}{\partial t} + \mathbf{v} \cdot \nabla f \right) d\mathbf{v} = \int_{\mathbf{v}} \mu(\mathbf{v}) (S_c + S_v(f)) d\mathbf{v}, \quad (2)$$

where $\mu(\mathbf{v}) = 1$ for the continuity equation, $\mu(\mathbf{v}) = m\mathbf{v}$ for the momentum equation (the velocity component parallel to the magnetic field is often the most important), and $\mu(\mathbf{v}) = m|\mathbf{v}|^2/2$ for the energy equation, with m the particle mass and $\int_{\mathbf{v}} \dots d\mathbf{v}$ the integral over the whole velocity space. When assuming a fluid distribution based on a Chapman–Enskog expansion [25] for f , we obtain a closed set of equations. For plasma edge neutrals, it is typically a valid approach to reduce the 3D momentum equation to a single momentum equation parallel to the magnetic field, to capture the dominant ion–neutral friction along that direction. For the perpendicular direction and in the absence of plasma drifts, pressure diffusion approximations suffice [26], obtained by neglecting perpendicular inertia terms in the neutral momentum equation. These simplifications are based on a scale-analysis of the terms assuming perpendicular velocities are small compared to the parallel and thermal velocities [26]. Extensions to the context of drifts that become of the order of these velocities are subject of ongoing research. In this project, we have significantly improved the existing fluid models for neutral atoms by deriving the transport coefficients and boundary conditions from the underlying kinetic description (AMJUEL/HYDHEL databases for the collisional processes and the TRIM database for the reflection physics, assuming an ideal flat surface) [3, 16, 27]. Additionally, an approximate treatment of the

effect of molecules is possible through modified boundary conditions. In this way, the transport model and boundary conditions for the fluid neutrals are made consistent with those of the full kinetic simulation, using the same tabulated data, and without introducing ad hoc tuning parameters. In the remainder of the paper, we call these improved models *AFNs*. Usually, one assumes equal ion and neutral temperatures for a fluid treatment, but reference [28] shows that solving a separate neutral energy equation further reduces the fluid–kinetic discrepancies. With the AFN models, excellent qualitative and quantitative agreement between fluid and kinetic simulations has been achieved for HCRs, without the need for ad hoc parameter tuning in our improved fluid neutral models, both for standalone neutral and coupled neutral–plasma simulations. Moreover, the accurate numerical discretization of the (isotropic) neutral transport fluxes on grids misaligned with respect to the magnetic field was proven essential in this picture [29]. Through the use of boundary conditions based on TRIM data, the impact of wall materials on the reflection properties is now captured by the fluid models [16, 30]—an effect that was inaccessible with earlier fluid neutral approaches.

A pure fluid treatment of the neutral atoms introduces a modelling error that may be quite significant, especially in lower recycling regimes outside of the strict validity range of the models. However, there are obvious benefits in terms of computational cost, with fluid neutral simulations typically being at least an order of magnitude faster than kinetic simulations. The reduced accuracy may be acceptable for initial scoping studies or large parameter scans. If the error is not acceptable, the hybrid methods discussed below can provide an answer.

1.2. Spatial hybridization (SpH)

The validity of the fluid model is restricted to regions of high (mainly CX) collisionality, where the \mathbf{K}_n is sufficiently small. This typically occurs only in a small part of a dense divertor. In other regions, in particular the main chamber, the private flux region (PFR), and other remote areas such as sub-divertor structures, kinetic effects need to be incorporated. Various flavours of *SpH approaches* have been proposed in the literature, usually distinguishing between kinetic and fluid domains, with proper boundary conditions at the interfaces. The exact treatment of the boundary or interface conditions depends on the specific flavour of the hybrid method. Intuitively, the interface has to be placed in a region that is ‘fluid enough’ so that boundary conditions can be defined based on fluid-like distributions, but not ‘too far’ into the fluid so that the computational gain is large. For some methods the interface has to be specified by the user, for others the method does it automatically based on criteria of, e.g., local collisionality to determine evaporation/condensation reactions. In the context of the linear Boltzmann equation, this domain decomposition can be approached in a different way, relying on the localisation of physical sources. Indeed, in the *SpH* presented in [20], we exploit the linearity of the kinetic equation to split the source term into separate contributions, deciding for each whether to treat the neutrals as kinetic or fluid: $\sum_i S_{c,i} =$

$\sum_i S_{c,i}^k + \sum_i S_{c,i}^f$, leading to a corresponding decomposition of the solution $\sum_i f_i \approx \sum_i f_i^k + f^f$, where the superscript now denotes the treatment, and all sources treated through a fluid model are included in the equation for a single fluid population f^f :

$$\frac{\partial f_i^k}{\partial t} + \mathbf{v} \cdot \nabla f_i^k = S_{c,i}^k + S_v(f_i^k), \quad (3)$$

$$\frac{\partial f^f}{\partial t} + \mathbf{v} \cdot \nabla f^f = \sum_i S_{c,i}^f + S_v(f^f). \quad (4)$$

There are no mutual interactions between the fluid and kinetic parts of the solution. Note that this approach would have to be further developed to account for neutral–neutral collisions, which makes the problem nonlinear. Including these collisions can then be achieved by resolving the additional nonlinearity introduced by the BGK model across iterations, as is already the standard procedure for BGK collisions in EIRENE at the moment, but will not be discussed further here. It is immediately obvious that for the source terms $S_{c,i}^k$ treated kinetically, no modelling error is introduced on the corresponding part of the distribution f_i^k . This typically concerns source terms due to recycling at main chamber and PFR boundaries, as well as gas puff sources. The sources due to recycling of ions at the targets, as well as those arising from volumetric recombination, are grouped into equation (4), which is then approximated with the AFN models described above, equation (2), typically solved as part of the plasma fluid solver (meaning additional input to EIRENE from the CFD side of the package). At plasma–vacuum boundaries, fluid neutrals can continue their trajectory as kinetic through appropriate interface conditions [20]. The approximation of equation (4) with a fluid model introduces a modelling error on this part of the population. However, in the SpH approach the fluid population is only ‘fed’ with sources in HCRs, where the accuracy of the fluid models is high. The resulting fluid neutral density far from the HCRs is negligible. As a result, the modelling error in the SpH approach is limited, while a significant speed-up is achieved. At present the decision to treat a certain source as fluid or kinetic is still up to the user; methods to automate the decision based on local K_n criteria are being developed.

The SpH method for atoms can easily be coupled to a standard kinetic model for molecules. Thermally emitted particles at a surface are typically launched as (D_2) molecules, while fast reflected particles are treated as (D) atoms [1]. For fluid atoms impinging on a surface, an estimate of the thermally reflected fraction follows directly from the boundary conditions of the AFN. Upon dissociation, the molecules give rise to another source contribution $S_{c,i}$ of atoms, for which it has to be decided whether to treat them (i) kinetically (entering equation (3)); or (ii) as a fluid by adding them as a source term in the fluid neutral model (through equation (4)). These options were explored in [31]. Method (i) is more accurate, whereas method (ii) leads to the largest speed-up. For JET L-mode cases, simulations with the SpH neutral model are on average between 7 and 14 times faster than simulations with fully kinetic neutrals. In that study, the number of MC particles used in the kinetic and hybrid methods was not the

same, but was carefully chosen to have the same statistical error on the relevant output quantities of interest, thereby allowing an honest assessment of the speed-up. The hybrid method needs much fewer particles to achieve a given statistical error, since neutrals originating from some of the sources are treated as a fluid (which does not itself introduce additional statistical noise).

In the approach discussed so far, the whole domain is either fluid or kinetic for a given source. In the **SpH** implemented in [21], with main idea following [32], for all atom sources the distribution function is split into a kinetic and fluid part in the whole domain, so that for a given source fluid and kinetic neutrals co-exist. Additional reaction channels are introduced to represent evaporation/condensation between the kinetic and fluid phases, $S_{f \rightarrow k}$ and $S_{k \rightarrow f}$, respectively:

$$\frac{\partial f^k}{\partial t} + \mathbf{v} \cdot \nabla f^k = S_c + S_v(f^k) + S_{f \rightarrow k} - S_{k \rightarrow f}, \quad (5)$$

$$\frac{\partial f^f}{\partial t} + \mathbf{v} \cdot \nabla f^f = S_v(f^f) - S_{f \rightarrow k} + S_{k \rightarrow f}. \quad (6)$$

By construction, the sum of equations (5) and (6) gives the original kinetic equation (1). Moreover, equation (6) is then approximated through the solution of a fluid model, introducing a modelling error that may also affect the kinetic part of the solution through the coupling terms. At birth, all neutrals are treated kinetically (source S_c in equation (5)). The rate coefficients for evaporation/condensation depend on K_n and are chosen such that the trajectories of kinetic atoms are quickly stopped when entering a fluid region and contribute to the source of fluid atoms there. For the detailed transition criteria, we refer to reference [21]. The expected computational gain comes from this shortening of the trajectories in HCRs. The implementation of these models is straightforward and does not require modifications of the kinetic solver. However, the transition between fluid and kinetic regions/boundaries is based on ad hoc criteria, whose choice affects the accuracy of the scheme. If the K_n dependence of the condensation/evaporation rates is sharp enough so that the transition between domains where the fluid (resp. kinetic) phase strongly dominates occurs over a few grid cells, this method can be seen as an immersed boundary implementation of a standard SpH technique. Such a sharp transition ensures the ‘quick’ termination of kinetic particle trajectories when entering into fluid regions thereby optimizing the speed-up.

As future work, we plan to combine both SpH methods, by investigating optimal combinations of source splitting and evaporation/condensation reactions. For both SpH methods, the kinetic parts of the distribution f_i^k are positive by definition and represent a real part of the distribution for the total neutral population. The corresponding neutrals are simulated with a regular MC procedure, requiring limited implementation effort. The disadvantage of the SpH methods is that they are not completely equivalent with the kinetic description, but introduce a model error. The first SpH method is only accurate if the fluid limit is reached near the sources treated as fluid. The accuracy of the second SpH method strongly depends on the spatial distribution of the condensation and evaporation

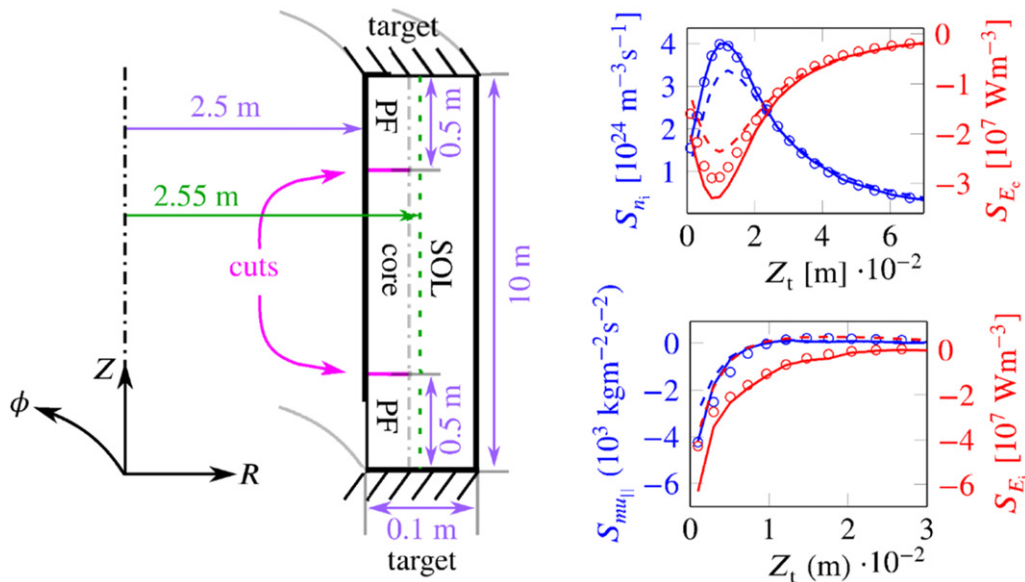


Figure 2. An example of FKH simulation for a slab test case, JET-relevant by the characteristic sizes and plasma parameters, using SOLPS-ITER (including EIRENE-NGM): particle, momentum and energy sources estimated with kinetic (solid lines), fluid (dashed line) and mMH (circles) simulations for a flux tube at 2.55 m. One can see that hybrid, in this case mMH, solution approaches the full kinetic one. [34] John Wiley & Sons. © 2020 WILEY-VCH Verlag GmbH & Co. KGaA, Weinheim.

sources. These disadvantages disappear when using a mMH approach as introduced in the next section.

1.3. Micro–macro hybrid (mMH) approach

In the *mMH approach* of reference [22], the (complete) neutral distribution function f is split into a fluid part f^f and a ‘kinetic correction’ f^δ seamlessly in the entire simulation domain, such that $f = f^f + f^\delta$ exactly holds. A consistent set of equations is derived in such a way that the sum of the fluid part and the kinetic correction gives exactly the same solution as the original fully kinetic equation. For the fluid part of the distribution, this leads to the AFN equations, with additional corrections on source and transport terms due to the kinetic correction population. The kinetic correction is obtained from a modified kinetic simulation involving positive and negative correction particles, with net zero density and momentum and provides an exact closure for the AFN model. As a result, the method requires more substantial development efforts for the kinetic correction terms, and requires a fluid solver acting for the whole volume right down to the first wall [33]. The benefit of mMH is that the solution is equivalent to the solution of the fully kinetic equation and independent of the recycling regime, illustrated in figure 2 for a JET-like slab case. However, for a significant computational benefit, the distribution implied by the fluid model should already be close to the fully kinetic distribution (hence, the kinetic correction part should be small), which is expected only for high-recycling or detached conditions. In regions of the simulation domain where kinetic effects are important, substantial cancellation errors arise, that presently hamper the efficiency and convergence of the method [34]. Note that the efficiency (speed-up) for individual source estimation on fixed plasma background is typically highest for the particle source, and lower for the higher order velocity moments (momentum/energy sources) which are more

strongly impacted by kinetic effects [34]. When looking at the coupled plasma–neutral simulations, this will lead to a single ‘effective’ speed-up over the entire simulation. To alleviate these effects, further research will look into projection techniques to limit the cancellation errors [22], as well as a synergistic combination with the SpH approach where the mMH technique is not applied to the complete neutral distribution, but only to the part of the distribution treated as fluid in the SpH approach. In this case, $f = \sum_i f_i = \sum_i f_i^k + f^f + f^\delta$ is again an exact decomposition. The terms f_i^k that show predominantly kinetic behaviour are treated with a standard kinetic simulation, whereas the mMH approach will provide an exact closure for the part of the neutral population treated as fluid.

Although both are MC simulations, the nature of a kinetic correction simulation needed by the mMH approach is substantially different from an MC simulation of the ‘regular’ kinetic problem of equation (1). Because the largest share of the CX collisions is captured by the fluid model, absorption events are much more dominant in the kinetic correction equation, while scattering (CX) events dominate the regular kinetic problem. A rigorous analysis of optimal combinations of MC tracking and source estimation schemes demonstrates that non-analog track-length or next-event tracking schemes are optimal under such conditions [15] while the default treatment in EIRENE is an analog particle tracking procedure. Further speed-up of the mMH approach can be expected by using these optimal tracking procedures [34]. The insights from [15] may also serve to optimize regular kinetic simulations by automatically selecting the best combination of tracking and estimator schemes based on local collisionality.

1.4. Kinetic–diffusion Monte Carlo method (KDMC)

The mMH and SpH approaches both require an additional discretization of the spatial domain suitable for the fluid

description (usually as part of the fluid plasma solver) and intricate couplings between fluid and kinetic neutral parts. An alternative that avoids these complications is the *kinetic-diffusion MC scheme (KDMC)*, an *asymptotic-preserving* scheme which is fully MC (APMC) [23]. In this method, the MC particles themselves behave alternately according to the kinetic equation and according to the fluid limit. During kinetic phases, the MC particles are simulated with a standard particle tracing method. During the fluid phases, the MC particles move via a random walk method, an MC discretization of the limiting fluid equation. In the KDMC-scheme, a random walk step is used after each collision during the kinetic phase. This random walk step continues until the end of the time step, the size of which is the numerical parameter of the scheme. The coefficients of the random walk step are designed to such that the positional increments match the kinetic process's mean and variance exactly in every time step. Furthermore, the correlation of the motion between subsequent time steps is maintained by the proposed combination of kinetic and fluid-like parts. Outside of the truly kinetic or truly fluid regimes, the KDMC scheme results in a bias. This small bias in the simulation determined by the time step is removed with a newly developed multilevel extension of the algorithm, where the multiple time step sizes define the levels [15, 35]. The final challenge for this method is the design of estimators for the random walk steps that capitalize on the achieved speed-up. Such estimators have been developed for a 1D simulation, and are being extended to higher dimensions.

2. Kinetic ions

In the previous section, several approaches designed to treat neutrals at least partially as a fluid in order to reduce the computational cost have been presented. Here we discuss the opposite situation, in which ion species are treated partially in a kinetic approach within the EIRENE solver. This is relevant for e.g. molecular ions involved in plasma edge chemistry, but also potentially for lowly ionized impurity ions. In fact, the relaxation time for the velocity distribution of the latter towards a local Maxwellian (through Coulomb collisions) can be longer than their ionization time. In such cases, where features of the neutral impurity velocity distribution can propagate to ions, only a kinetic approach can capture these effects, which should impact the parallel transport force balance for impurities. Therefore, a kinetic approach for plasma ions at trace levels is developed inside the EIRENE code. In fact, treating ions on the kinetic side of CFD-EIRENE packages in addition to a fluid/hybrid approximation for neutrals provides more flexible and seamless coupling between the codes as well as an internal benchmarking mechanism. In the last two years fundamental physical enhancements of the kinetic ion transport part of EIRENE were performed [36], by adding first-order drift effects, cross-field diffusion, and magnetic mirror force. These additions, which are relevant for thoroughly investigating the full three-dimensional influence of impurities on actual fusion devices, have been cross-checked on simple model cases against analytical properties of passing and trapped (banana) particle orbits, as well as checking on the

introduction of numerical diffusion by our integration scheme. These features are still under development in preparation for use in full physics cases.

3. Parallelisation of EIRENE

The MC method implemented in EIRENE follows independent particles (in between iterations with a CFD code), even in non-linear cases involving collisions between tracked species, where the code relies on an iterative BGK approach [19]. Massive parallelisation is thus in principle straightforward, as it involves communications only before and after the MC calculation itself. For a given accuracy (statistical noise level) parallelisation will reduce the wall clock time, or for a given wall clock time the accuracy can be improved. Reducing the statistical noise by one order of magnitude requires 100 times more test particles because of the square-root scaling characteristic of MC approaches. Because of the stratified sampling used in EIRENE (where one stratum often corresponds to a physical source, see section 1), different strategies are possible: the strata can be treated sequentially with particles distributed on all cores for each strata; or the strata can be treated in parallel, each one being allocated a number of cores on the basis of predefined criteria (involving the number of particles and the relative flux associated to a strata). The latter reduces the amount of communications but load balancing must be carefully considered. Both strategies were available in EIRENE and rely on a distributed memory parallelization using the MPI library. However, for large 3D grids used for transport calculations (e.g. EMC3-EIRENE) on large machines or for turbulent calculations on present day devices the memory requirements to store the ~ 100 input and output tallies of EIRENE become large enough so as to preclude running with one MPI process per core on a node. Moreover, when coupled to hybrid OpenMP/MPI parallelized codes such as SolEdge3X [11] the current implementation leads to an inefficient CPU resource usage when running coupled to EIRENE and running with more than one thread per core. As a result, an OpenMP shared memory layer has been added into EIRENE, where all the particles attributed to an MPI process are distributed among threads on the local node. This allows efficient use with hybrid OpenMP/MPI codes and has the potential to reduce the memory consumption by a factor of at most $1/N_{\text{core}}$, where N_{core} is the number of cores on the node. All tallies are treated as shared arrays within threads for a given MPI process and are therefore not duplicated in memory. In practice, the actual reduction in memory usage comes close to this limit as illustrated in figure 3(b), where using 48 cores implies only $\sim 10\%$ increase in memory while the same run with 48 MPI process implies a ~ 48 fold increase. However, figure 3(a) shows that the parallel efficiency strongly decreases above 12 cores on this architecture (Marconi-fusion [37]), because of the need of ensuring correctness when updating shared tallies (using ATOMIC clauses), but also because of memory access bottlenecks. Indeed, the current structuring of tallies in memory allows strong flexibility but is not optimal w.r.t. to OpenMP (or as well for usage of GPU-based systems which are typically also with shared memory), and addressing these shortcomings

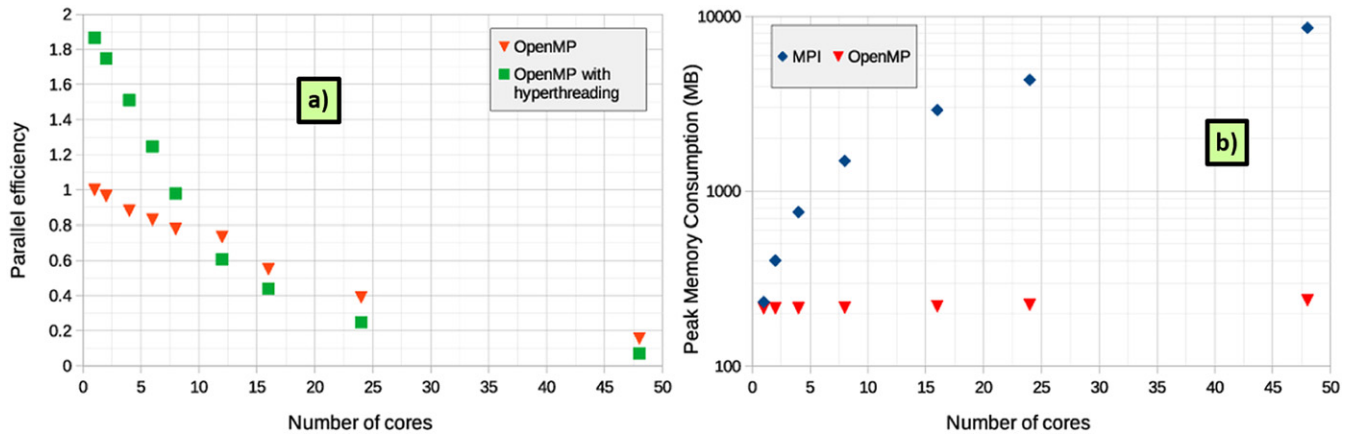


Figure 3. OpenMP/MPI hybrid parallelisation of EIRENE-NGM demonstrated on the ITER test case. The OpenMP efficiency (left, (a)) diminishes with the number of cores/threads, however the drop is monotonic and consistent. On a positive side OpenMP provides dramatic reduce of the peak memory consumption (right, (b)), which is critical for large and detailed simulation cases in particular in a view of the upcoming ITER and DEMO predictive modelling.

would require a major code restructuring. On the longer term, further relaxing memory constraints will require implementing a domain decomposition approach, with nodes being allocated chunks of the full grid, and particles distributed among threads on the node.

4. Discussion and perspectives of FKH approaches

A summary of the main advantages and remaining issues of the various fluid and hybrid methods is presented in table 1. Figure 4. gives an overview for the hybridisation methods discussed in this paper, in terms of the balance between computational speed (increases towards the left of the figure) versus modelling accuracy (increases towards the right), which are the main performance criteria for hybrid kinetic–fluid approaches, and typically need to be balanced against each other. To find this balance, one needs to investigate how both the accuracy and computational cost of each of the methods scales with the method parameters. Then, one needs to decide on either computational budget and choose those method parameters that maximise accuracy, or, conversely, decide on a desired accuracy and choose those method parameters that minimise the computational cost. It is quite natural to expect the conclusion to depend on both the case studied and the desired accuracy (or the available computational budget). Hence, a detailed comparison of different hybrid schemes that leads to meaningful conclusions is a difficult challenge. Here, we restrict ourselves to indicating the factors that influence the comparison, and overview a systematic way of optimising method parameters for each of the hybrid methods.

Computational errors in the fluid–kinetic context generally consist of discretization errors, cancellation errors, statistical errors, biases, and modelling errors. Except for modelling errors, these errors can be lowered by adjusting the method parameters, hereby increasing the computational cost. For instance, for the discretization errors, this is done by refining the mesh. For the statistical errors and biases, increasing the number of samples reduces the error [38].

AFN, SpH, and mMh cope with discretization errors for the respective fluid parts. To optimally choose grid sizes, guidelines are discussed in [38], on a coupled plasma–neutral level. The guidelines provided there also apply for the neutral simulation on its own: by using a sequence of coarsened/refined grids, the size of the discretization error can be estimated via Richardson extrapolation. Additional errors present in the particle (MC) part of SpH and in mMh and KDMC that have to be taken into account are the errors arising due to stochasticity. In a simple coupled plasma–neutral simulation, there is on the one hand a statistical error that scales as $1/\sqrt{N}$ with N the number of particles. On the other hand, a bias due to the nonlinear coupling with the plasma simulation is present that scales as $1/N$. Simulation strategies with averaging over different runs, results in different scalings, as shown and discussed at length in [39].

In [38] a methodology is presented that jointly selects the grid size and the number of particles optimally for a given computational time. The modelling error present in AFN and SpH for instance, depends critically on which application case is considered and, in the case of SpH, also on how the boundary between the submethods is defined. The truly hybrid methods such as SpH, mMh, and KDMC benefit from a large scale separation between different regions of the application case, meaning that there are some clearly kinetic regions and some clearly fluid regions.

All methods have difficulty with regions of intermediate, often called rarefied, behaviour. The precise balance struck between scale separation, accuracy, and code performance, is a key point of further investigation. In this text, we will only go over the key effects. For SpH, this intermediate region has to be located and the boundary between its two submethods has to be established there. If the boundary is taken close to the truly fluid region, the resulting model error will be low, but at the cost of a sharp increase in computational cost, and vice versa if the boundary is close to the truly kinetic region. The mMh method will automatically result in a fully accurate simulation of the intermediate region, but if the scale separation is low, this might result in very poor code performance in

Table 1. An overview of implemented FKH approaches.

FKH approach	Main idea and parameter(s)	Advantages	Issues	Development status and performance gain
Advanced fluid neutral models (AFN) [16]	Replace kinetic simulation with a fluid model, tailored to conditions of high CX collisionality	<ul style="list-style-type: none"> - Large speed-up compared to kinetic simulation (> order of magnitude) - No statistical noise - Tight coupling with plasma equations may improve convergence 	<ul style="list-style-type: none"> - Introduces a modelling error, that may be substantial in conditions/regions - Requires a dedicated grid/solver (typically as part of the plasma code) 	<ul style="list-style-type: none"> - Implemented in SOLPS-ITER - Good accuracy for high recycling conditions, incl. ITER, demonstrated [27, 30] - Speed-up factor > 10
Spatial (SpH)-based on source location [31]	Treat neutrals born in high K_n regions as kinetic, and born in low K_n regions as fluid	<ul style="list-style-type: none"> - Straightforward implementation and coupling to molecules - Clear improvement in accuracy compared to AFN 	<ul style="list-style-type: none"> - (Small) remaining modelling error - Choice of kinetic/fluid source treatment up to user - Sub-optimal speedup because kinetic trajectories entering HCR not terminated 	<ul style="list-style-type: none"> - Implemented in SOLPS-ITER - Good accuracy for high recycling conditions, incl. ITER, demonstrated - Speed-up factor ~10 demonstrated for JET L-mode discharges, incl. molecules [31]
SpH with evaporation / condensation [21]	Co-existence and interaction of two phases (kinetic and fluid) on the full domain, can be made	More seamless transition, automatic procedure for spatial	-additional assumptions e.g. CX dominating more complicated than fixed source	<ul style="list-style-type: none"> - Implemented in SOLedge2D EIRENE - potential for speed up demonstrated (30% reduction of EIRENE CPU time in moderately collisional cases) - porting to SOLedge3X and combination with Spatial SpH are in elaboration
micro-Macro (mMH) [22]	Based on exact decomposition of kinetic equation in fluid and kinetic correction parts	Modelling error can be completely removed, at all collisionalities	Substantial <ul style="list-style-type: none"> -Implemented in low collisional /kinetic regions hamper convergence - Requires complete overlap of grids for fluid and kinetic correction neutrals grid (up to the wall) 	<ul style="list-style-type: none"> -Implemented in SOLPS-ITER -Speed-up factor ~5-10 demonstrated in simplified geometries
APMC option: Kinetic diffusion (KDMC) [23] multilevel approach (ML-KDMC [35])	Particles follow a hybridized path that combines advection diffusion and kinetic steps	<ul style="list-style-type: none"> -No need to resolve individual collisions -No need for separate fluid/neutral grid - free of cancellation - bias can be suppressed by ML-KDMC 	Currently only available for singlespecies scattering/absorption	The approach is for now just demonstrated on simplified problems, treatment of simulation cases similar to e.g. SOLPS-ITER tasks is in elaboration.

some cases. The impact of rarefied regions on KDMC depends on the interplay with the used time discretization. If the time step used is sufficiently small, also the rarefied region will be accurately simulated. This comes at a computational cost that is (at most) inversely linear to the time step size. For the

KDMC method, a solution to this issue has been developed by incorporating the KDMC method into the multilevel MC framework [35], which combines simulations with different time steps. The cheap, large time-step simulations allow to quickly reduce the statistical error. The expensive, small

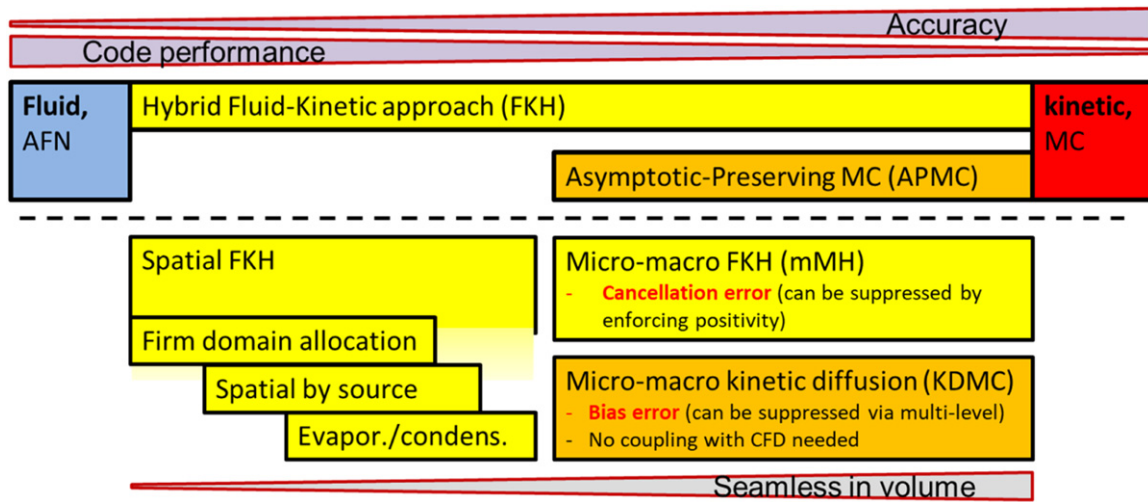


Figure 4. Hierarchy of neutral modules (above the dashed line) and hybrid fluid–kinetic approaches already implemented in EIRENE (below the dashed line).

Table 2. Alternative ways to increase performance.

Approach	Main idea and parameter(s)	Combinability with FKH?	Development status and performance gain
Parallelisation (brute force approach)	<ul style="list-style-type: none"> - Strong scaling allows reducing the wall clock time or improving the MC statistics - Weak scaling allows more detail and larger simulation volume. neutrals) 	Provide additional performance on the remaining kinetic calculation (largest effect expected on APMC where no fluid solver is required for	<ul style="list-style-type: none"> - Good MPI scaling - Decent OpenMP efficiency within 10-20 threads/cores - About 100 times reduction of the memory peak consumption.
Optimised tracking of kinetic processes (selection of critical states, bundled states etc.)	<p>Selection of simulation/estimator procedure based on analysis of variance</p> <p>Improve CRM performance by reorganisation and scalings for e.g. isotopologically resolved hydrogen data</p>	<ul style="list-style-type: none"> - Combination with mMH and SpH straightforward - Combination with APMC (e.g. KDMC) requires new source term estimators <p>The same reaction rates can be utilized by both the fluid and kinetic parts of FKH codes</p>	<p>Results in prototype code, can be transferred to EIRENE</p> <ul style="list-style-type: none"> - Schemes for CRM construction exist. - Isotopologically resolved hydrogen data is still scarce

time-step simulations allow to recover the full kinetic accuracy.

Besides the main criteria of accuracy and code performance, there are several more qualitative aspects that differentiate the methods. Methods consisting of multiple methods such as SpH and mMH are more difficult to implement than KDMC and AFN, but also the incorporation of those methods into EIRENE may require significant changes to the existing code. The purely particle-based method of KDMC, is furthermore trivially parallelizable. Other methods have the advantage of being more mature in their development such as SpH and AFN. In conclusion, the decision on which of the many existing methods is the best choice in the long term is still open. In this section, we stress the need for sound choices

in grid size and particle numbers as in [38] and we identify the error-accuracy balance in rarefied regimes for SpH, mMH, and KDMC as an important research topic.

As discussed above, further research will investigate synergistic combinations of the various hybrid approaches, aiming for at least an order of magnitude speed-up for the most accurate hybrid schemes. To our view, we need to keep two FKH tracks under development: one based on AFN/SpH/mMH and another based on the asymptotic-preserving KDMC scheme. The first track should evolve to a single unified option combining the two SpH options via the use of the condensation/evaporation terms, and providing exact closure for the fluid neutral population through an mMH approach. This line is already strongly developed, however further development

is needed specifically to limit the cancellation error in the mMH scheme. KDMC accuracy is compromised in a similar way by a bias error at intermediate collisionalities, which however can be completely suppressed by the multilevel ML-KDMC. In the end, the final performance and flexibility on DEMO scale will be decisive. In addition, the first (AFN) track will always enable tighter coupling with the plasma equations (could be beneficial for convergence of the coupled system) and should have lower MC noise (beneficial for sensitivity analysis as used e.g. in gradient-based optimization studies [40, 41]). The KDMC approach, on the other hand, allows standalone use of EIRENE without iterations with the CFD code, which is beneficial for geometry flexibility and following of A & M processes including kinetic ions.

Table 2 shows alternative ways to increase performance, which should work independent on the simulation type—hybrid or full kinetic except to optimisations. One can expect also some performance gain from general restructuring and refactoring of the code, including the use of the modern, HPC-friendly output format HDF5, providing e.g. parallel writing by multiple cores. Also the accurate computation of sensitivities in the presence of MC noise (adjoint approach) has been demonstrated [43], and will be pursued in further EIRENE development involving algorithmic differentiation as shown in [41] for the (deterministic) plasma solver within SOLPS-ITER.

5. Summary and conclusion

FKH approaches are developed (for instance SpH and mMH gaining performance due to a partial use of the AFN approximation [16]) for the CFD-EIRENE packages [20, 21, 34]. They combine improved computing performance with model accuracy approaching full kinetic simulations. The alternative perspective APMC approach is also considered [23], including development and first tests of the newly proposed KDMC formulation with a multilevel option (ML) [35] aimed to overcome the bias. In addition, the option to track ions kinetically is improved [36]. The advantages of hybridisation methods are compared based on experience from the first applications to test cases relevant for fusion devices. Currently, the main effort is on (1) basic development of the approaches (2) validation with full-kinetic simulations to determine the gain in computational speedup and optimal parameters (3) impact demonstration of new physics included on, for example, ITER-relevant applications (4) unification of the methods allowing e.g. combined mMH and SpH simulations as indicated in this paper. The methodology [38] is formulated allowing enveloping the error-accuracy balance in rarefied regimes for various FKH approaches as well as uniform guidelines for selecting the optimal grid size and the number of particles. In the end we aim to obtain a NGM providing a hierarchy of neural (and trace ion) models with certain guidelines and automatic tools for hybridisation parameter optimisation allowing a user to select an appropriate FKH option according to his accuracy/performance requirements.

At the present stage, we recommend to continue development of 2 FKH lines in parallel: the merged into

one AFN/SpH/mMH approach (the way for that is suggested) and the ML-KDMC approach [23]. Each of the approaches is expected to have an automatic accuracy control and self-adaptation mechanism for any case at hand.

The hybrid OpenMP–MPI code parallelization and optimisation of the A & M process treatment (improved CRMs) go mostly in parallel adding an additional factor to the improvement of the EIRENE-NGM performance. However, this factor can depend on the final selection of the FKH scheme and overall optimisation of the code.

Acknowledgments

The paper presents a joint effort within the EUROfusion Theory and Advanced Simulation Coordination (E-TASC), task ‘NGM’ 2019–2020 and the related O-EIRENE HLST project. 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. W. Van Uytven is funded by a PhD fellowship of the Research Foundation Flanders (FWO). Parts of the work were supported by the Research Foundation Flanders (FWO) under project Grant G078316N.

ORCID iDs

D.V. Borodin  <https://orcid.org/0000-0001-8354-1387>
 F. Schluck  <https://orcid.org/0000-0003-4634-0679>
 S. Wiesen  <https://orcid.org/0000-0002-3696-5475>
 S. Brezinsek  <https://orcid.org/0000-0002-7213-3326>
 W. Dekeyser  <https://orcid.org/0000-0003-1954-3729>
 S. Carli  <https://orcid.org/0000-0001-6581-2745>
 M. Blommaert  <https://orcid.org/0000-0002-2616-8042>
 W. Van Uytven  <https://orcid.org/0000-0002-7253-2087>
 M. Baelmans  <https://orcid.org/0000-0001-9905-4583>
 B. Mortier  <https://orcid.org/0000-0003-1788-7819>
 G. Samaey  <https://orcid.org/0000-0001-8433-4523>
 Y. Marandet  <https://orcid.org/0000-0002-9970-538X>
 E. Westerhof  <https://orcid.org/0000-0002-0749-9399>
 J. Gonzalez  <https://orcid.org/0000-0001-7905-5001>
 M. Groth  <https://orcid.org/0000-0001-7397-1586>
 A. Holm  <https://orcid.org/0000-0002-8391-8364>

References

- [1] Reiter D., Baelmans M. and Börner P. 2005 The EIRENE and B2-EIRENE codes *Fusion Sci. Technol.* **47** 172–86
- [2] Summers H.P. 2004 *The ADAS User Manual, Version 2.6* (<http://adas.ac.uk>)
- [3] Reiter D., Janev R.K. *et al* 2002 EIRENE-related data: AMJUEL, HYDKIN, HYDHEL, H2VIBR *Berichte des Forschungszentrums Jülich: Juel-3966, Juel-4005, Juel-4038, Juel-4105* (http://eirene.de/html/a-m_data.html)
- [4] Fantz U., Reiter D., Heger B. and Coster D. 2001 Hydrogen molecules in the divertor of ASDEX Upgrade *J. Nucl. Mater.* **290–293** 367–73

- [5] Groth M. *et al* 2011 Influence of cross-field drifts and chemical sputtering on simulations of divertor particle and heat loads in ohmic and L-mode plasmas in DIII-D, AUG, and JET using UEDGE *J. Nucl. Mater.* **415** S530–4
- [6] Wiesen S. *et al* 2015 The new SOLPS-ITER code package *J. Nucl. Mater.* **463** 480–4
- [7] Bonnin X., Dekeyser W., Pitts R., Coster D., Voskoboynikov S. and Wiesen S. 2016 Presentation of the new SOLPS-ITER code package for tokamak plasma edge modelling *Plasma Fusion Res.* **11** 1403102
- [8] Simonini R. *et al* 1994 Models and numerics in the multi-fluid 2D edge plasma code EDGE2D/U *Contrib. Plasma Phys.* **34** 368–73
- [9] Bufferand H. *et al* 2015 Numerical modelling for divertor design of the WEST device with a focus on plasma–wall interactions *Nucl. Fusion* **55** 053025
- [10] Frerichs H., Bonnin X., Feng Y., Loarte A., Pitts R. A., Reiter D. and Schmitz O. 2019 Stabilization of EMC3-EIRENE for detachment conditions and comparison to SOLPS-ITER *Nucl. Mater. Energy* **18** 62–6
- [11] Bufferand H. *et al* 2021 Progress in edge plasma turbulence modelling—hierarchy of models from 2D transport application to 3D fluid simulations in realistic tokamak geometry *Nucl. Fusion* **61** 116052
- [12] Pitts R.A. *et al* 2019 Physics basis for the first ITER tungsten divertor *Nucl. Mater. Energy* **20** 100696
- [13] Stangeby P. 2000 *The Plasma Boundary of Magnetic Fusion Devices* (Bristol: Institute of Physics Publishing) p 180
- [14] Schwartz J.A., Emdee E.D., Goldston R.J. and Jaworski M.A. 2019 Physics design for a lithium vapor box divertor experiment on magnum PSI *Nucl. Mater. Energy* **18** 350–5
- [15] Mortier B. 2020 Advanced Monte Carlo simulation and estimation for kinetic neutral particles in the plasma edge of fusion reactors *Dissertation PhD Thesis* KU Leuven, Department of Computer Science
- [16] Horsten N., Samaey G. and Baelmans M. 2017 Development and assessment of 2D fluid neutral models that include atomic databases and a microscopic reflection model *Nucl. Fusion* **57** 116043
- [17] Fleck J.A. Jr and Canfield E.H. 1984 A random walk procedure for improving the computational efficiency of the implicit Monte Carlo method for nonlinear radiation transport *J. Comput. Phys.* **54** 508–23
- [18] Börgers C., Larsen E.W. and Adams M.L. 1992 The asymptotic diffusion limit of a linear discontinuous discretization of a two-dimensional linear transport equation *J. Comput. Phys.* **98** 285–300
- [19] Crestetto A., Crouseilles N. and Lemou M. 2012 Kinetic/fluid micro–macro numerical schemes for Vlasov–Poisson–BGK equation using particles *Kinetic Relat. Models* **5** 787
- [20] Blommaert M., Horsten N., Börner P. and Dekeyser W. 2019 A spatially hybrid fluid–kinetic neutral model for SOLPS-ITER plasma edge simulations *Nucl. Mater. Energy* **19** 28–33
- [21] Valentinuzzi M., Marandet Y., Bufferand H., Ciraolo G. and Tamain P. 2019 Two-phases hybrid model for neutrals *Nucl. Mater. Energy* **18** 41–5
- [22] Horsten N., Samaey G. and Baelmans M. 2020 A hybrid fluid–kinetic model for hydrogenic atoms in the plasma edge of tokamaks based on a micro-macro decomposition of the kinetic equation *J. Comput. Phys.* **409** 109308
- [23] Mortier B., Baelmans M. and Samaey G. 2020 Kinetic-diffusion asymptotic-preserving Monte Carlo algorithms for plasma edge neutral simulation *Contrib. Plasma Phys.* **60** e201900134
- [24] Romazanov J. *et al* 2019 Beryllium global erosion and deposition at JET-ILW simulated with ERO2.0 *Nucl. Mater. Energy* **18** 331–8
- [25] Bobylev A.V. 1982 The Chapman–Enskog and Grad methods for solving the Boltzmann equation *Dokl. Akad. Nauk SSSR* **262** 71–5
- [26] Schneider R. *et al* 2006 Plasma edge physics with B2-Eirene *Contrib. Plasma Phys.* **46** 3–191
- [27] Blommaert M., Dekeyser W., Horsten N., Boerner P. and Baelmans M. 2018 Implementation of a consistent fluid–neutral model in SOLPS-ITER and benchmark with EIRENE *Contrib. Plasma Phys.* **58** 718–24
- [28] Van Uytven W., Blommaert M., Dekeyser W., Horsten N. and Baelmans M. 2020 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
- [29] Dekeyser W., Bonnin X., Ligo S. W., Pitts R. A. and LaBombard B. 2019 Implementation of a 9-point stencil in SOLPS-ITER and implications for Alcator C-Mod divertor plasma simulations *Nucl. Mater. Energy* **18** 125–30
- [30] Van Uytven W. *et al* 2022 Assessment of advanced fluid neutral models for the neutral atoms in the plasma edge and application in ITER geometry *Nucl. Fusion* **62** 086023
- [31] Horsten N., Groth M., Blommaert M., Dekeyser W., Pérez I.P. and Wiesen S. 2021 Application of spatially hybrid fluid–kinetic neutral model on JET L-mode plasmas *Nucl. Mater. Energy* **27** 100969
- [32] Karney C.F.F., Stotler D.P. and Braams B.J. 1998 Modeling of neutral plasma in a divertor in the fluid–kinetic transition *Contrib. Plasma Phys.* **38** 319–24
- [33] Dekeyser W. *et al* 2021 Plasma edge simulations including realistic wall geometry with SOLPS-ITER *Nucl. Mater. Energy* **27** 100999
- [34] Horsten N., Dekeyser W., Blommaert M., Samaey G. and Baelmans M. 2020 A hybrid fluid–kinetic neutral model based on a micro-macro decomposition in the SOLPS-ITER plasma edge code suite *Contrib. Plasma Phys.* **60** e201900132
- [35] Mortier B., Robbe P., Baelmans M. and Samaey G. 2022 Multilevel asymptotic-preserving Monte Carlo for kinetic-diffusive particle simulations of the Boltzmann-BGK equation *J. Comput. Phys.* **450** 110736
- [36] Schluck F. 2020 Integration of kinetic ions in a three-dimensional Monte-Carlo neutraltransport code *Contrib. Plasma Phys.* **60** e201900144
- [37] Iannone F. *et al* 2018 MARCONI-FUSION: the new high performance computing facility for European nuclear fusion modelling *Fusion Eng. Des.* **129** 354–8
- [38] Ghoo K., Börner P., Dekeyser W., Kukushkin A.S. and Baelmans M. 2018 Grid resolution study for B2-EIRENE simulation of partially detached ITER divertor plasma *Nucl. Fusion* **59** 026001
- [39] Ghoo K., Dekeyser W., Samaey G., Börner P. and Baelmans M. 2016 Accuracy and convergence of coupled finite-volume/Monte Carlo codes for plasma edge simulations of nuclear fusion reactors *J. Comput. Phys.* **322** 162–82
- [40] Baelmans M., Blommaert M., Dekeyser W. and Van Oevelen T. 2017 Achievements and challenges in automated parameter, shape and topology optimization for divertor design *Nucl. Fusion* **57** 036022
- [41] Carli S., Blommaert M., Dekeyser W. and Baelmans M. 2018 Sensitivity analysis of plasma edge code parameters through algorithmic differentiation *Nucl. Mater. Energy* **18** 6–11
- [42] Holm A., Börner P., Rognlien T.D., Meyer W.H.W.H. and Groth M. 2021 Comparison of a collisional-radiative fluid model of H₂ in UEDGE to the kinetic neutral code EIRENE *Nucl. Mater. Energy* **27** 100982
- [43] Dekeyser W., Blommaert M., Ghoo K., Horsten N., Boerner P., Samaey G. and Baelmans M. 2018 Divertor design through adjoint approaches and efficient code simulation strategies *Contrib. Plasma Phys.* **58** 643–51