First time-dependent SOLPS-ITER simulations of detachment burn-through on MAST-U

D Moulton[1], R. Scannell[1], J. Flanagan[1], S. Leigh[1], the MAST-U team[a] and the EUROfusion Tokamak Exploitation team[b]

[1]UKAEA, Culham Campus, Abingdon OX14 3DB, UK

[a]See author list of J. Harrison et al 2024 (https://doi.org/10.1088/1741-4326/ad6011)

[b]See author list of N. Vianello et al 2026 Nucl. Fusion in press (https://doi.org/10.1088/1741-4326/ae71ec)

Introduction

We present the first interpretative simulations of transient burn-through on MAST-U in a Super-X divertor (SXD) configuration, using fully time-dependent SOLPS-ITER. Simulations agree well with measured IR peak heat loads and Fulcher emission for Type-I ELMs. We find that simulations do not align with a model whereby a single energy threshold for burn-through exists, set by the energy required to ionise the pre-existing neutral reservoir. Instead, a time threshold is observed in the simulations that corresponds to the parallel convection time, followed by a loading of the target that is strongly dependent on target recycling (no recycling leads to a factor ~5 higher peak target heat loads compared to instantaneous recycling).

Comparison to analytic solution for 1D isothermal plasma expansion into vacuum

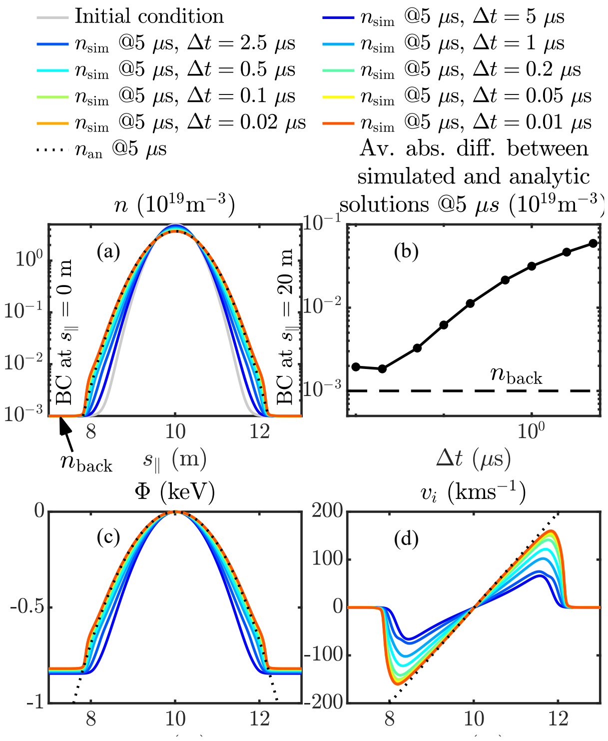

Figure 1.(a,c,d) Simulated values (solid lines) tend to analytic expectation (dotted lines) as time step is reduced (blue to red). (b) Difference between simulated and analytic solutions saturates around the background density $n_{\rm back}$ required for stable numerics.

Before applying the code to a realistic geometry with time-dependent Eirene, the standalone B2.5 time-dependent model (instigated with the input switch 'b2mndt_style' set to 2) was verified against the analytic solution for a 1D isothermal expansion of a gaussian density perturbation into vacuum [Mora 2005]. With sufficient spatial and time resolution, good agreement was found between the analytic and simulated profiles of density, potential and parallel velocity (Figure 1a,c,d). A small background density $n_{\rm back}$ was necessarily added for stable solutions and the difference between analytic and simulated density tended towards $n_{\rm back}$ with decreasing time step Δt. It's important to note that one must use the latest code version (since 20/3/2026) in order to obtain this correct behaviour (otherwise timescales were found to scale inversely with the number of internal iterations due to a bug introduced on 9/2/2023). Comparison to the expressions for adiabatic expansion (allowing time variation of electron temperature T_e) given in [Mora 2005] were not possible since his energy equation assumed a constant T_e. A less vigorous verification of the time-dependent Eirene was also performed via comparison to an in-house neutral code and experimental dry runs (not shown).

Elm simulation setup

A pure deuterium version of the SXD simulation shown in [Verhaegh 2025] was reconverged in time-dependent mode (b2mndt_style=2, Eirene time census activated with a time horizon equal to the B2.5 time step of $\Delta t = 10^{-5}$ s, warm restart [NPTST=0] with 5×10^7 histories, plus a further 2000 new histories launched from each strata at each new step). This time-dependent background simulation had an input power of 1 MW and, for this initial study, we used the same L-mode transport coefficients as in [Verhaegh 2025]. The resulting neutral pressure, shown in Figure 2a, was ~1 Pa in the sub-divertor. The simulation required 750 GB memory over 36 cores and was found to match well to its equivalent quasi-steady-state simulation.

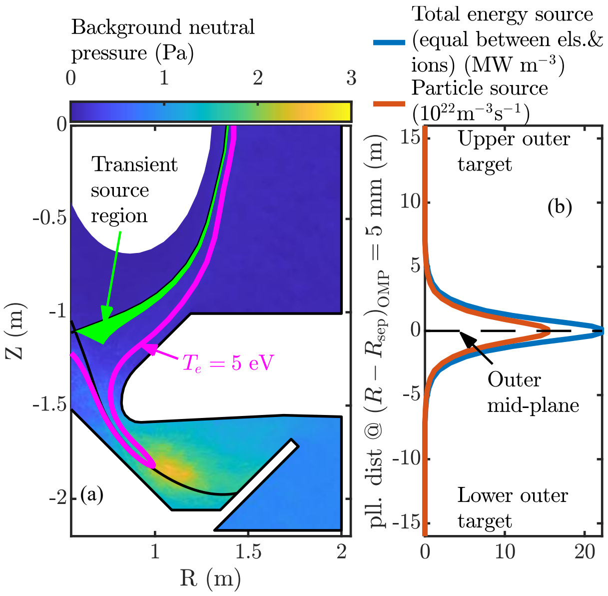

Figure 2. Background neutral pressure and T_e =5 eV front used as an initial condition before imposing the source shown in (b) for a duration of 1 ms.

In this initial scoping study, we use this same background throughout and impose an instantaneous upstream source inside the green region in Figure 2a for a duration of 1 ms. The source was radially constant over a width λ_{trans} =1 cm and poloidally varying as shown in Figure 2b. The plasma energy (S_E) and particle (S_n) sources were set such that.$S_E = 3k_B T_{\mathrm{ped}} S_n$ [Havlickova 2012], with T_{ped} =300 eV chosen here and electrons and ions given equal energy in the reference cases. It was a deliberate simplification to have an instantaneous source since this removes the additional timescale of the source itself from our analysis. In what follows, the total transient power P_{trans} (i.e. the space integral of S_E) was scanned from 1 MW to 10 MW (corresponding to a scan in transient energy ΔW from 1 kJ to 10 kJ).

Comparison to experiment

Simulations were compared to Ultra Fast Divertor Spectroscopy (UFDS) Fulcher emission (see Figure 3a for lines of sight drawn over the simulated Fulcher emission for the background plasma) and Infra-red (IR) target heat loads of ELMs on MAST-U. Experimental data (black dots in Figure 3b and c) are taken from

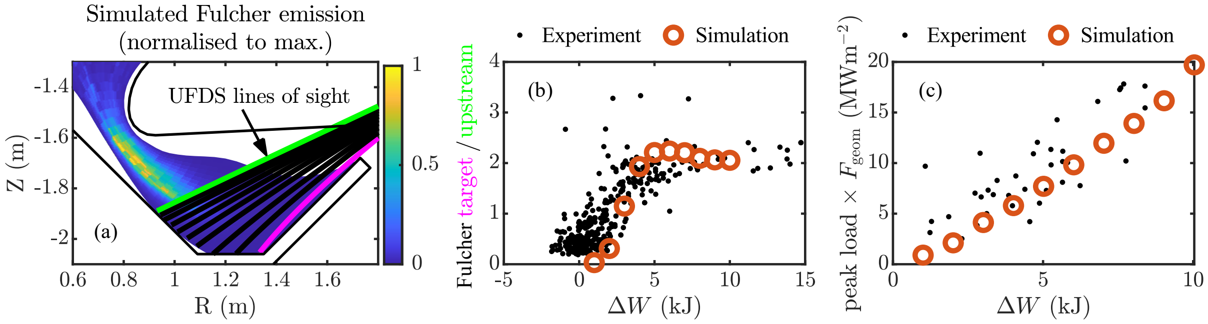

Figure 3. (a) UFDS lines of sight over simulated Fulcher emission for the background simulation. (b) Mean target/upstream line of sight values (coloured as in (a)) over the transient duration, in simulation and experiment. (c) Peak target load (multiplied by F_{geom}.

[Flanagan 2025] and [Flanagan 2026], respectively (see those references for details). In general, the agreement between simulation and experiment is acceptable, especially considering that no effort was made to tune these simulations. In figure Figure 3b, the mean of the line-integrated target-divided-by-upstream Fulcher emission is compared (same data as figure 7e from [Flanagan 2025]). In Figure 3c, F_{geom} is the geometry factor defined in [Scannell 2026] which allows different geometries to be compared fairly (the geometry simulated had more poloidal flux expansion than the experimental data shown).

The concept of a burn-through 'threshold'

Figure 4a shows the simulated peak (lower outer) target load as a function of time for P_{trans} =1 MW to P_{trans} =10 MW, as labelled. There is clearly a *time* threshold τ_{delay} that the transient source needs to exist for, so that the plasma has time to reach (and load) the target. The notion of having to 'burn-through' a pre-existing neutral reservoir might lead one to expect $\tau_{\mathrm{delay}} \sim W_{\mathrm{neut}}/P_{\mathrm{trans}}$, where $W_{\mathrm{neut}} = N_{\mathrm{neut}}\varepsilon_{iz}$ (N_{neut} and ε_{iz} are the total number of neutrals in the pre-existing neutral reservoir inside λ_{trans} and the average cost to ionise a neutral, respectively).

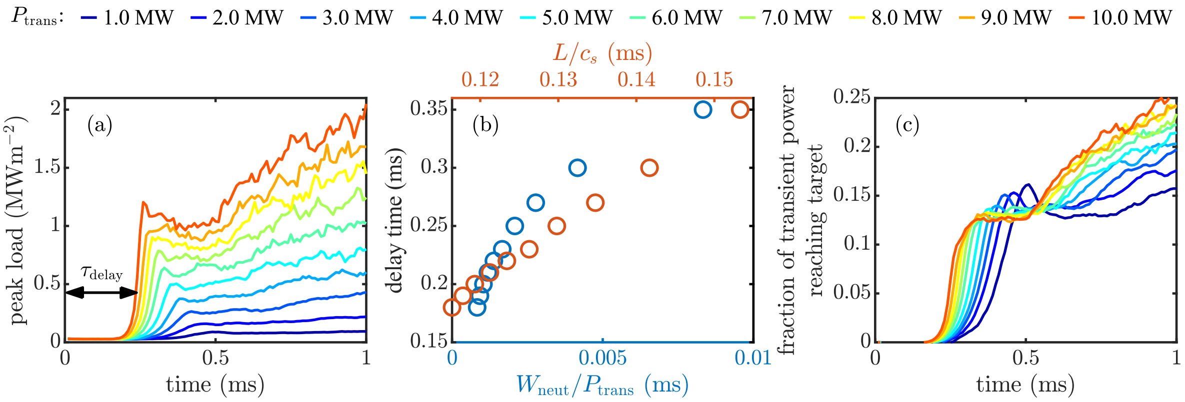

Figure 4.(a) Target loading as a function of time for the reference simulations. (b) Comparison of τ_{delay} as a function of L/c_s (red) and $W_{\mathrm{neut}}/P_{\mathrm{trans}}$ (blue). (c) fraction of transient power reaching the target as a function of time.

However, we find that τ_{delay} does *not* scale linearly with $W_{\mathrm{neut}}/P_{\mathrm{trans}}$ but *does* with the parallel convection timescale $\tau_{\parallel} = L/c_s$ (averaged at the outer divertor entrance). This is demonstrated in Figure 4b, which compares the two scalings. Absolute values of $W_{\mathrm{neut}}/P_{\mathrm{trans}}$ are also far too small to explain the delay time since W_{neut} is just 12.4 J with ε_{iz} =30 eV. Thus, the energy required to ionise the pre-existing neutral reservoir is insufficient to stop target loading. Once the target is reached, recycling starts to dictate the loads that are reached (see below). The fraction of transient power reaching the target after 1 ms is similar for all powers (Figure 4c). Thus, the final load scales approximately linearly with ΔW (open circles in figure Figure 3c).

The neutral reservoir and the importance of recycling

By default, we assume instantaneous recycling of the incoming target ion flux (solid lines in Figure 6). In that case, for a ΔW =5 kJ transient whose 5 eV front just reaches the target after 1 ms, the neutral reservoir inside λ_{trans} is only depleted by ≲30% before being boosted again by recycling. By contrast, when the target is fully pumped at the instant the transient is initialised (dashed lines), the neutral reservoir is completely depleted and the peak load to the target is ~5 times higher. This implies that neutrals do play an important role in reducing peak heat loads, but neutrals born from recycling rather than the pre-existing reservoir (this point was also previously made by [Smith 2020]). In the default case, the 5 eV front also slows once the incoming ion flux starts to recycle from the target, which is not observed in the absence of recycling.

Increased ion heating share leads to increased loads

There is uncertainty in our choice of how the transient energy source is shared between electrons and ions. In high powered simulations (ΔW =10 kJ), a 3 × bias towards the ions lead to significantly higher (factor $\gtrsim$2) peak target loads when the transient first reaches the target (Figure 7). The time profile of the loading is reminiscent of the case without recycling shown in Figure 6b. This suggests that recycled neutrals are less able to dissipate incoming ion heat compared to electron heat. We hypothesise that this could be a possible mechanism for 'high burn-through events', where increased peak target heat loads occur for the same ΔW for some sawtooth transients [Scannell 2026].

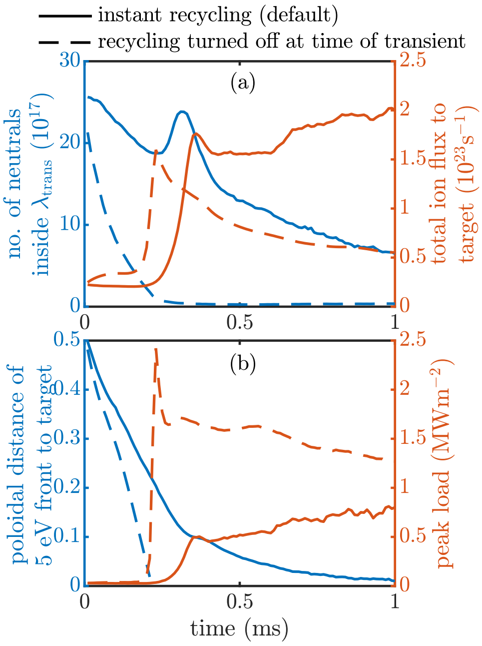

Figure 5. Impact of turning off recycling at the instant the transient is launched (dashed lines), compared to the default (instantaneous recycling) case (solid lines)

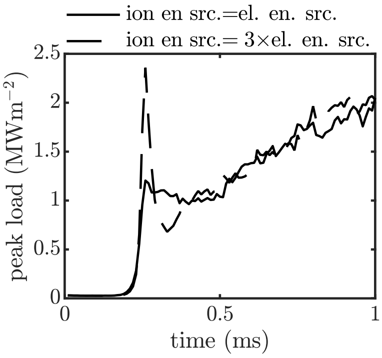

Figure 6. Impact of biasing the transient power source towards the ions by a factor 3 compared to the electrons.

Future work

This paper comprises just an initial scoping study. We intend to compare loading in Super-X vs Conventional configurations in order to assess whether the Super-X offers benefit beyond geometry factors alone (experimental evidence suggests that this depends on the type of transient and on whether there is nitrogen seeding or not [Flanagan 2026, Scannell 2026]). We will repeat the presented transient scan for a range of different divertor neutral pressures. We will also assess timescales for detachment (rather than attachment, as studied here), once the transient source is removed, and compare to experiment. Finally, we intend to include impurities (both intrinsic carbon and seeded nitrogen), although we note that cases with impurities quickly become highly memory intensive for time-dependent Eirene.

References

Flanagan J et al (2025) Nucl. Fusion **65** 116031
Flanagan J et al (2026) PSI conference, Regensburg, Germany
Havlickova E et al (2005) Plasma Phys. Control. Fusion **54** 045002
Mora P (2005) Phys. Plasmas **12** 112102
Scannell R et al (2026) IAEA-FEC conference, submitted to Nucl. Fusion
Scannell R et al (2026) This conference
Smith SF et al (2020) NF **60** 066021
Verhaegh K et al (2025) Commun. Phys. **8** 215

Acknowledgements

This work has 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) and from the EPSRC Fusion Grant [grant number EP/W006839/1]. 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.