UK Atomic Energy Authority

UKAEA-STEP-PR(26)01

M. Khan, M. Kryjak, I. Cziegler, B. Dudson, S. Henderson, D. Moulton, R. Osawa, C. Ridgers

The role of momentum transfer in the detachment front response to power transients for reactor scale tokamaks

The role of momentum transfer in the detachment front response to power transients for reactor scale tokamaks

M. Khan, M. Kryjak, I. Cziegler, B. Dudson, S. Henderson, D. Moulton, R. Osawa, C. Ridgers

This is a preprint of a paper submitted for publication in Nuclear Fusion

The role of momentum transfer in the detachment front response to power transients for reactor scale tokamaks

Matthew Khan

School of Physics, Engineering and Technology, University of York, YO10 5DQ, United Kingdom

E-mail: matthew.khan@stfc.ac.uk
Central Laser Facility, STFC Rutherford Appleton Laboratory, Harwell Oxford, United Kingdom

Mike Kryjak

School of Physics, Engineering and Technology, University of York, YO10 5DQ, United Kingdom
UKAEA (United Kingdom Atomic Energy Authority), Culham Campus, Abingdon, Oxfordshire, OX14 3DB, UK

István Cziegler

School of Physics, Engineering and Technology, University of York, YO10 5DQ, United Kingdom

Benjamin Dudson

Lawrence Livermore National Laboratory, 7000 East Avenue, Livermore CA 94550, USA

Stuart Henderson

UKAEA (United Kingdom Atomic Energy Authority), Culham Campus, Abingdon, Oxfordshire, OX14 3DB, UK

David Moulton

UKAEA (United Kingdom Atomic Energy Authority), Culham Campus, Abingdon, Oxfordshire, OX14 3DB, UK

Ryoko Osawa

UKAEA (United Kingdom Atomic Energy Authority), Culham Campus, Abingdon, Oxfordshire, OX14 3DB, UK

Christopher P. Ridgers

School of Physics, Engineering and Technology, University of York, YO10 5DQ, United Kingdom

May 2024

Abstract. We have investigated the response of detached plasmas to multi-ms timescale power transients using 1D simulations of the scrape-off layer on a scale and parameter range relevant for reactor-scale tokamaks, such as the Spherical Tokamak for Energy Production (STEP). Simulations were performed using the fluid code Hermes-3, employing the solver CVODE for robust time integration. Including or disregarding the momentum received by the neutral gas via charge exchange in simulations is found to dramatically affect the detachment front's response. In the analysis, motion of the ionisation front in a detached plasma is determined by two distinct mechanisms; a fast initial pressure-dominated response as the neutral cloud is rapidly compressed, and a slow burn-through dominated one as the cloud is ionised. The former is only observed if the momentum transfer to neutrals is kept. With transfer of momentum to the neutrals disabled, the plasma only interacts with the neutrals through ionisation and recombination. Thus, the front speed during the transient is determined by the rate of burn-through of the neutrals only. Results are compared to a recent semi-empirical reattachment model [S.S. Henderson et al. 2024 Nucl. Fusion 64 066006], which shows agreement over a range of power transient forms and powers, where uncertainties are encompassed by the range of potential fits provided by model parameter choice.

Keywords: Divertor, Detachment, Scrape-off layer, Power Transient, STEP, time-dependent simulation

Submitted to: *Plasma Phys. Control. Fusion*

1. Introduction

The Spherical Tokamak for Energy Production (STEP) [1, 2, 3, 4] is a reactor-scale device currently in the design phase that aims to produce 100 MWe net electrical power. A key requirement of these devices is the ability to manage high-power loads transported from the plasma core across the separatrix and into the scrape-off layer (SOL). Rapid plasma transport parallel to the magnetic field brings heat, mass, and momentum to plasma-facing components in the SOL. Without proper mitigation techniques, the heat flux experienced by these plasma-facing components would be beyond the capabilities of any material, leading to damaging effects such as melting and sputtering [5]. With the addition of high atomic number (Z) impurities, radiation within the SOL can reduce the plasma temperature to the point where recombination can occur and a gas cloud forms at the divertor wall (also referred to as the target) and the plasma is said to detach. The detachment front, defined in this work as the point where the plasma temperature drops to 5 eV, separates the SOL plasma from the neutral cloud. The front and neutral gas dissipate energy and momentum primarily through radiation and charge-exchange reactions, respectively. Thus, detachment is essential for the longevity of plasma-facing components and must be maintained throughout high-performance operation, which is especially challenged by transient fluctuations in power entering the SOL. These fluctuations might be caused by intentional actions such as tokamak start-up [6], the L-to-H transition [7], fuel pellet injection [8], power sharing variation in double-null divertor configurations [9, 10], or from uncontrolled events such as the eruption of ELM filaments [11]. Investigations into how detached plasmas respond to transients are essential for understanding how they might resist the destruction of the detached state, i.e. reattachment.

Large-scale integrated SOL simulations, such as SOLPS-ITER [12], are often employed to describe divertors and include sophisticated physics models for both plasmas and neutrals. However, the computational demands of these simulations are substantial, making them less suitable for solving dynamic situations such as transients, which involve both short and long timescales [13], and especially for the high throughput needed for optimisation studies. Models such as DLS (Detachment Location Sensitivity) [14, 15, 16] can estimate the response of the detachment front to a change in upstream conditions, but are not time-resolved as needed for transient events. Though recent work has developed a reduced complexity semi-empirical model for the reattachment timescale [17], this model is thought to be incompatible with pulses of large enough magnitude to cause complete reattachment, and fast (tens of ms and below) transient phenomena have not been considered.

More complete (though still somewhat reduced) models, such as 1D fluid simulations, are well-suited for performing explorative studies to complement the large-scale simulations. By employing reasonable assumptions, the 3D geometry of a divertor can be reduced to a single 1D flux tube, greatly reducing computation time. This opens the possibility for fast and wide-reaching investigations into detachment response to transients.

In this paper, we present a 1D simulation study of the response of a detached plasma in a STEP-like reactor to power transients. The aim of this study was to investigate the behaviour of the system in regimes dominated by either momentum loss or ionisation (burn-through). The difference between these two regimes is the inclusion of a key term in the fluid evolution equations, F_{CX}, that controls the transfer of momentum from the plasma to the neutral species via charge exchange reactions. Comparisons of simulations neglecting and including this term are used to highlight the importance of charge exchange momentum transfer in determining the response of the detachment front to the transient. This work is performed via 1D simulations using the Hermes-3 fluid code [18, 19], built using the BOUT++ framework [20]. High-order time integration is provided by the CVODE solver [21], which is used to solve fluid equations for density and pressure of ions, electrons, and neutrals, as well as for the momentum of ions and neutrals.

The rest of this paper is organised as follows. Section 2 introduces the Hermes-3 model. Section 3 motivates the choices for the simulation parameters and the shapes of the power transient pulses modelled. Results are given in Section 4, first for steady-state solutions. The semi-empirical model is then compared with the Hermes-3 simulations where the momentum-transfer term in the neutral fluid evolution equations is disabled (no-F_{CX} case) and enabled (full case). This is done for rapid rises in power, then slow rises, and finally, finite-length power pulses (i.e. a rapid rise followed by a plateau and then a rapid drop in the power).

2. Hermes-3 Model

Hermes-3 is a multi-dimensional fluid code built using the BOUT++ framework. In this work a total of seven spatially varying quantities are evolved in time using a conservative finite differences scheme; deuterium ion and atom densities (n_{d+} and n_d), flow of ions and atoms parallel to the magnetic field (v_{d+}, v_d) and the pressure of the ions, atoms, and electrons (p_{d+}, p_d, and p_e). Quasineutrality and a dominant hydrogenic species are assumed such that $n_e = n_{d+}$, and in this work the parallel current is set to zero such that $v_e = v_{d+}$.

The following 1D fluid equations are solved for the deuterium ions.

$$\frac{\partial n_{\mathrm{d}+}}{\partial t} = -\,\partial_\parallel\left(n_{\mathrm{d}+}v_{\mathrm{d}+}\right) + S_{PI} + S_{iz/rec} \tag{1}$$

$$\begin{aligned}\frac{\partial}{\partial t}\left(m_{d+}n_{d+}v_{\parallel,d+}\right) =& -\,\partial_\parallel\left(m_{d+}n_{d+}v_{d+}\cdot v_{d+}\right) - \partial_\parallel p_{d+} \\ &+ en_{d+}E_\parallel + F_{CX/iz/rec/collisions}\end{aligned} \tag{2}$$

$$\begin{aligned}\frac{\partial}{\partial t}\left(\frac{3}{2}p_{d+}\right) =& v_{d+}\partial_\parallel p_{d+} - \partial_\parallel\left[\frac{5}{2}p_{d+}v_{d+} - \kappa_{d+}\partial_\parallel T_{d+}\right] \\ &+ W_{input} + W_{CX/iz/rec/collisions} + R_{hyd/imp}\end{aligned} \tag{3}$$

Here, '$\parallel$' refers to the parallel direction along the magnetic field line, $\mathbf{B}$. κ_{d+} is the Braginskii parallel thermal conductivity. Electron force balance is used to calculate the parallel electric field $E_\parallel = (\mathbf{B}/B)\cdot\mathbf{E}$, which couples the electron and ion pressures.

Additional density sources, S, include an upstream density feedback controller that uses a proportional-integral (PI) controller to achieve a specified value, as well as sources from ionisation (iz) and recombination (rec). Frictional forces, F, are exerted on the ions from ionisation, recombination, charge exchange (CX), and collisions with neutrals. Pressure sources, W, are present for the reactions and collisions mentioned above, as well as the input power source from the core, which acts from the midplane to the X-point. R_{hyd} and R_{imp} represent the energy losses from hydrogenic (ionisation and recombination) and impurity radiation, respectively. Atomic reactions of ionisation, recombination and charge exchange are taken from the AMJUEL database [22] in the same fashion as in the Kotov model, as featured in SOLPS-ITER [23, 23], but omitting molecules. A fixed fraction model for impurities is used for this work, drawing the radiation rates from the ADAS database with an $n_e\tau$ value of $0.5\times10^{20}\,\mathrm{m}^{-3}\mathrm{ms}$, found to be in best agreement with simulations [24] of the ITER scrape-off layer.

The equivalent model for the electrons is much simpler than the model for deuterium ions in equations (1–3). This is due to the assumption of quasineutrality and zero current. Thus, only the electron pressure must be solved, which follows the same equation as that for the ions, with the radiation terms omitted.

The upstream boundary condition for the ions is reflective (no flux), such that $v_{d+} = 0$. This in effect assumes a double null configuration. The pressure and density employ a Neumann boundary condition for the upstream. Downstream, a Bohm sheath [25] is assumed such that the ion velocity is greater than or equal to the local sound speed, where the sheath extracts heat from both the ions and electrons. Ions that are reflected off the target return at a temperature of 3.5 eV, a typical energy resulting from molecular dissociation [26]. Density and temperature profiles are extrapolated to the target.

The neutral deuterium atomic species, d, obeys the following 1D fluid equations, similar to the equations for the ions,

$$\begin{aligned}\frac{\partial n_d}{\partial t} =& -\partial_{\|}\left(n_d v_d\right) - S_{iz/rc} \\ &+ \partial_{\|}\left(D_d n_d \frac{\partial_{\|} p_d}{p_d}\right)\end{aligned} \tag{4}$$

$$\begin{aligned}\frac{\partial}{\partial t}\left(m_d n_d v_{\|d}\right) =& -\partial_{\|}\left(m_d n_d v_d \cdot v_d\right) - \partial_{\|} p_d - F_{iz/rc/cx/collisions} \\ &+ \partial_{\|}\left(D_d m_d n_d v_d\left(\frac{\partial_{\|} p_d}{p_d} + \partial_{\|} T_d\right)\right)\end{aligned} \tag{5}$$

$$\begin{aligned}\frac{\partial \frac{3}{2} p_d}{\partial t} =& v_d \partial_{\|} p_d - \partial_{\|}\left[\frac{5}{2} p_d v_d - \kappa_d \partial_{\|} T_d\right] - W_{iz/rc/cx/collisions} \\ &+ \partial_{\|}\left(D_d \frac{3}{2} p_d \frac{\partial_{\|} p_d}{p_d}\right) + \partial \|\left(\frac{5}{2} D_d n_d \partial_{\|} T_d\right)\end{aligned} \tag{6}$$

where, $D_n = \left(\frac{B}{B_{pol}}\right)^2 \frac{T_n}{m_d \nu_d}$ is the diffusion factor, using the collision frequency ν_d and $\frac{B}{B_{pol}} = 10$ (projection of the cross-field direction). The diffusive terms are included to provide an approximation of neutral cross-field transport. Neutral particles use a no-flux boundary condition on both the upstream and downstream, such that $v_d = 0$. Profiles for density, temperature and pressure follow a Neumann boundary condition.

We will consider two particular cases. One where the above equations are solved in full, and another where the F_{CX} term (i.e. momentum transfer between ions and neutrals due to charge exchange) is kept for ion momentum equation but omitted for that of the neutrals. In doing so, F_{CX} will act as a loss term for the ions' momentum, but there will be no corresponding transfer of momentum to the neutrals and it is instead lost from the system. A comparison of these cases, referred to as 'full' and 'no-F_{CX}' respectively, will be used to show that charge exchange momentum transfer is crucially important in our simulations in the response of the detachment front to a transient power pulse.

3. Simulation Setup

The simulation configuration is based on an earlier version of the magnetic geometry and divertor setup of STEP and does not necessarily reflect the most recent design [27]. The focus of the present study is not to model the exact conditions of the STEP divertor, or to match simulations that use higher-fidelity models, but instead to explore the qualitative features of a system that is operationally relevant for STEP and to understand the principles of the plasma dynamics.

The parallel connection length of the simulated flux tube is 78.13 m, with a path length of 35.13 m from the midplane to the X-point and 45 m within the divertor from the X-point to the target. The flux expansion factor, the ratio of the flux tube area from the X-point to the target, is 2.051, where the full flux expansion profile along the magnetic field is taken from a magnetic equilibrium file. Note that rather than following the peak heat flux channel, a single flux tube is modelled. Recent 2D simulations have shown that the peak heat flux channel both widens and drifts radially when moving from the X point to the target [28]. For ITER, this effective flux expansion factor was shown to be ∼2× larger than the magnetic flux expansion for a single flux tube. This phenomenon is outside the scope of this study and will be addressed in future work. The computational grid contains 480 cells parallel to the magnetic field, which exponentially decrease in width towards the target. This reaches ≈mm cell sizes at the target and ≈10 cm cell sizes at 10 m from the target, which is a typical front detachment position used throughout this work. The field line used is $\lambda_q = 1.5\,\mathrm{mm}$ away from the separatrix, chosen to reflect the typical location of peak parallel heat flux [29], resulting in an unmitigated parallel heat flux at the X-point of $q_{||} \sim 770\,\mathrm{MWm^{-2}}$. An electron heat flux limiter of 0.2 was chosen, as this value was found to provide the best comparison against higher-fidelity kinetic models for ITER-like scenarios [30], though it should be noted that this value is not expected to be generally applicable [31, 32].

Throughout, a fixed fraction of the argon impurity of 1.5% and a target recycling fraction of 100% are used. The choice of complete recycling is to allow full control over the particle count in the system with other mechanisms. The deep detachment reached in this study results in extremely low target ion fluxes, effectively eliminating the impact of this value on the dynamics of the system once detached. Upstream densities of $\sim 4.4\times10^{19}\,\mathrm{m^{-3}}$ were found to produce parallel distances of detachment of ≈10 m from the target, approximately equal to half-way between the target and the X-point in poloidal space. The upstream density was set using a built-in proportional-integral (PI) controller, a feedback control system that achieves a user-specified value for the upstream density. It does so by applying a particle source or sink in the upstream region at a rate determined by the current difference of the upstream density from the set value (proportional term) and the historical difference (integral term). The gain factor on each of these terms can be altered to find the desired responsiveness of the PI controller as a whole, with higher gain factors resulting in a stronger and more immediate feedback. The gain values have been tuned to avoid controller overshoots

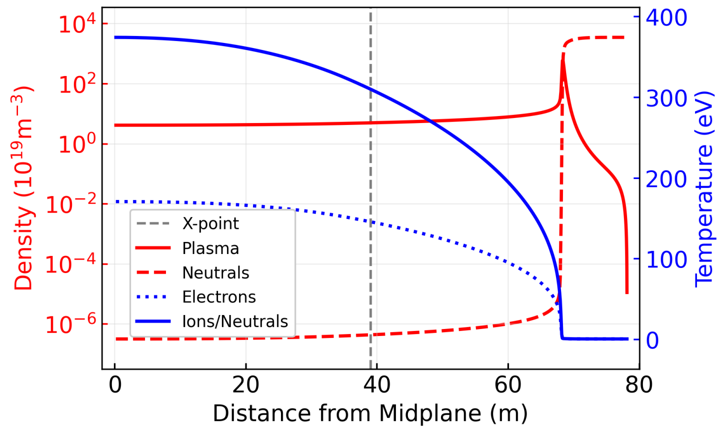

Figure 1. Density and temperature profiles against the parallel distance along the field line of a steady-state scenario for a STEP-like divertor. Electron and ion densities are identical, labelled as plasma. In deep detachments such as this, the neutral temperature relaxes to approximately that of the ion temperature. The vertical dashed line indicates the location of the X-point.

and to suppress oscillations.

Figure 1 shows the temperature and density profiles for a typical steady-state simulation. The upstream density was tuned to yield a detachment front position of $\sim 10\,\mathrm{m}$ from the target, defined here as the point at which $(T_e + T_{d+})/2 = 5\,\mathrm{eV}$. In the figure, this approximately corresponds to the peak in the plasma density interfacing with the neutral gas cloud, as well as the point at which the temperatures rapidly drop to near zero.

3.1. Transient Pulse

Transient events often result in heat, mass, or both being ejected from the confined plasma into the SOL. This study focuses on power transients only, which have no associated particle fluxes. During these events, the power crossing the separatrix into the simulated flux tube is increased from the baseline value to a chosen factor, where the simulation uses an equipartition of energy into electrons and ions. Additionally, there is no additional component for the impurities in the fixed fraction impurity model employed here. This increase in power is enacted through raising the energy density of the plasma components, i.e. through an increase in pressure.

Figure 2 shows the time history of a generalised model power transient. It features a hyperbolic tangent rise and fall in power. The rise time is defined as the duration it takes for the power to increase from 10% to 90% of the total power rise. The power pulse length is defined as the duration from the midpoint of the rise to the midpoint of the fall in power. However, a significant portion of the present study focuses on the initial transient phase, which involves a single tanh rise that is either very rapid, approximately resembling a step function within the system's timescales, or slow, similar to a ramp. The motivation is to determine the characteristic behaviour and timescales in response

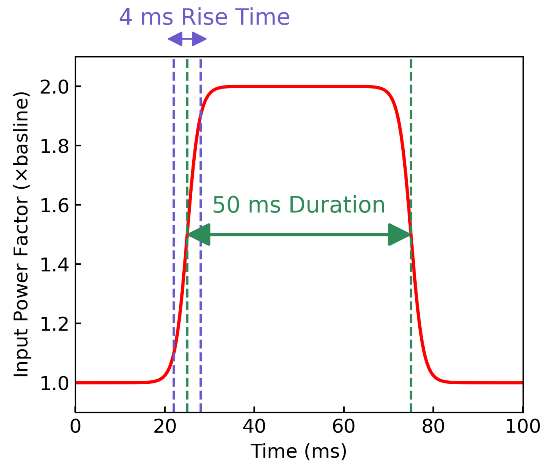

Figure 2. Example pulse used for power transients in this work, featuring a hyperbolic tangent profile for the power rise and fall. Thin blue dashed lines mark the 10% and 90% points in the rise, and thick green dashed lines show the midpoints of both the leading and trailing edges.

to a single power change, in the two aforementioned limiting cases.

An example of a phenomenon which can result in the rapid Heaviside-like transients mentioned above is a vertical displacement event in a tokamak, occurring via the natural vertical instability in diverted topology. The timescale of these is typically in the few ms range, and thus our choice of the fastest pulse at 4 ms approximately consistent with a single pressure wave equilibrating along the field corresponds well to such events. The most common concern in the case of double-null configurations, such as those designed for both MAST-U and STEP, is that the power load shared between the top and bottom divertors can be rapidly redirected to a single target [33] due to a movement in the plasma vertical position. Our choice therefore, is to explore the effects of a pulse with a magnitude in the range of $\sim 2\times$.

Slower power ramp transient cases are well-motivated by changes to the power crossing the separatrix, P_{SOL}, due to transitions in, for example, confinement and core heating, which reach the SOL via a plasma transport channel from the confined region, leading to a timescale on the energy confinement time τ_{E}. The latter was shown to be the case on a series of experiments performed on ASDEX-Upgrade and fitted well by the analytical model of re-attachment discussed later [17].

Finally, pulses of the type in Fig. 2, i.e. finite-duration transient pulses, are used to test for the presence and potential characteristic timescales of any hysteretic effect that might be observed in rapid transients which dissipate shortly after the rise. It is worth noting that the important limiting case of ELM filaments is not studied herein. The timescale of a single ELM filament transit is remarkably short, and perhaps more importantly, the geometry of the burst event is complicated by the high mode numbers and the narrow radial localisation of the filament. Consequently, a study of detachment response to an ELM is left for future work.

3.2. Consequences of Closed Particle System

A key feature that is not intrinsic to 1D simulations is a full description of cross-field transport, i.e. heat, mass, and momentum transfer from one flux tube to another. Without this phenomenon (and with the 100% recycling and neutral reflection boundary condition at the target), the particles are enclosed within the system and are unable to escape.

A consequence of a closed system is unphysically high resilience to reattachment. During the simulation of a power transient, heat is rapidly transported to the detachment front, where it increases the ionisation rate of the plasma-facing neutrals and the detachment front moves towards the target. The newly ionised plasma increases the local plasma density and flows down the density gradient seen in Figure 1 towards the midplane, increasing the upstream density as well. Without a loss of plasma particles from the flux tube via cross-field transport, the plasma density is seen to increase dramatically during a transient to values in excess of $5\times$ steady-state, far beyond what would be physically expected. This rise in plasma density across the flux tube increases radiative losses, reducing the heat flux that reaches the front, slowing the rate of further ionisation. This feedback loop results in the system resisting additional movement in the detachment front towards the target. Without modifications, extremely high powers, over $1000\times$ the steady-state power, were required to reattach the plasma, far exceeding what is observed in current devices and what is expected for STEP.

The increase in impurity radiation in our simulation is twofold. First, an increase in local electron density at the front will increase the radiative losses of the impurities, in line with expectations. The second source, however, is a consequence of the fixed fraction impurity model. Since the total impurity content is directly tied to the total number of plasma particles, the generation of new plasma through ionisation will result in new impurity particles being introduced to the system. In reality, the total number of impurity particles is fixed (ignoring the case of additional sources), and so their number would not increase with ionised neutrals.

In reality, the rise in density across a flux tube would increase cross-field transport, thereby limiting the peak values. To counteract this, the PI controller, which had typically only been used to establish a steady-state position, was kept enabled throughout the power transient to emulate cross-field transport. This eliminates the growth of plasma density above the set value, with the rate of removal proportional to the excess, just as large densities would lead to large cross-field fluxes. It was found that by implementing this approach, the reattachment powers and timescales were brought down to within the expected range. However, the PI controller is a crude stand-in for transport, and its influence on our results has been explored in Appendix B.

4. Results

in the simulations carried out, the movement of the detachment front was seen to behave via two distinct mechanisms. The first is a pressure-dominated movement where the increase in plasma pressure during a power transient is transferred to the neutrals, which then compress towards the target and the front moves in with it. The second, referred to here as burn-through, is the progression of the front due to the increased power delivered to the front that ionises the plasma-facing neutrals. As pressure is expected to equilibrate on a scale dictated by the sound speed, the pressure-dominated movement occurs on the same timescale as the change in input power, down to power rises of < 4 ms, and no residual front movement is seen once the input power reaches a constant value. Comparatively, burn-through is far slower, occurring on 10's to 100's of ms, and will occur throughout the transient event while the front is not in equilibrium, i.e. both pressure-dominated movement and burn-through will occur during the power change. As will become apparent later, the magnitude of the pressure driven movement can be much greater than that of burn-through resulting in the incorrect appearance that they are separated in time.

Within this present study, the controlling term in the equations given above was found to be the transfer of momentum between ions and neutrals in charge exchange collisions, F_{CX}. Only in the full cases is the pressure-dominated movement seen, whereas burn-through is seen in both the full case and no-F_{CX}. In all simulations, the F_{CX} term is included in the plasma momentum equations, resulting in momentum loss for the ions during charge exchange collisions. With it enabled for the neutrals, this momentum loss from the ions is transferred to the neutrals and, through this momentum exchange channel, the two species are said to be in pressure balance, such that pressure changes for one species exerts a force on the other. With the omittance of the term for the neutrals, the ion momentum is not transferred to the neutrals and is instead lost from the system. We revisit the physical consistency of this choice in Section 5.

In the no-F_{CX} cases, the neutrals are dominated by diffusion and respond very similarly to simulations that omit neutral momentum entirely. The evolution of neutral momentum is, at times, deliberately omitted from SOL models and simulations [28, 34], and thus disabling the F_{CX} provides a comparison against these setups. The distinct differences seen in the responses to transient events from the presence of the F_{CX} term highlights the need to consider it when modelling the SOL.

In this section, the impact of disabling F_{CX} is first shown through a comparison in steady-state, followed by the three types of transients studied (rapid rise, slow rise, and finite length pulses), and a comparison to the semi-empirical model.

4.1. Steady State Comparison

Figure 3 shows the steady state profiles for the two cases; full simulations and no-F_{CX}. The plasma density and pressure remain relatively unchanged with the inclusion of the F_{CX} term, however, the neutral properties can be seen to change dramatically. The

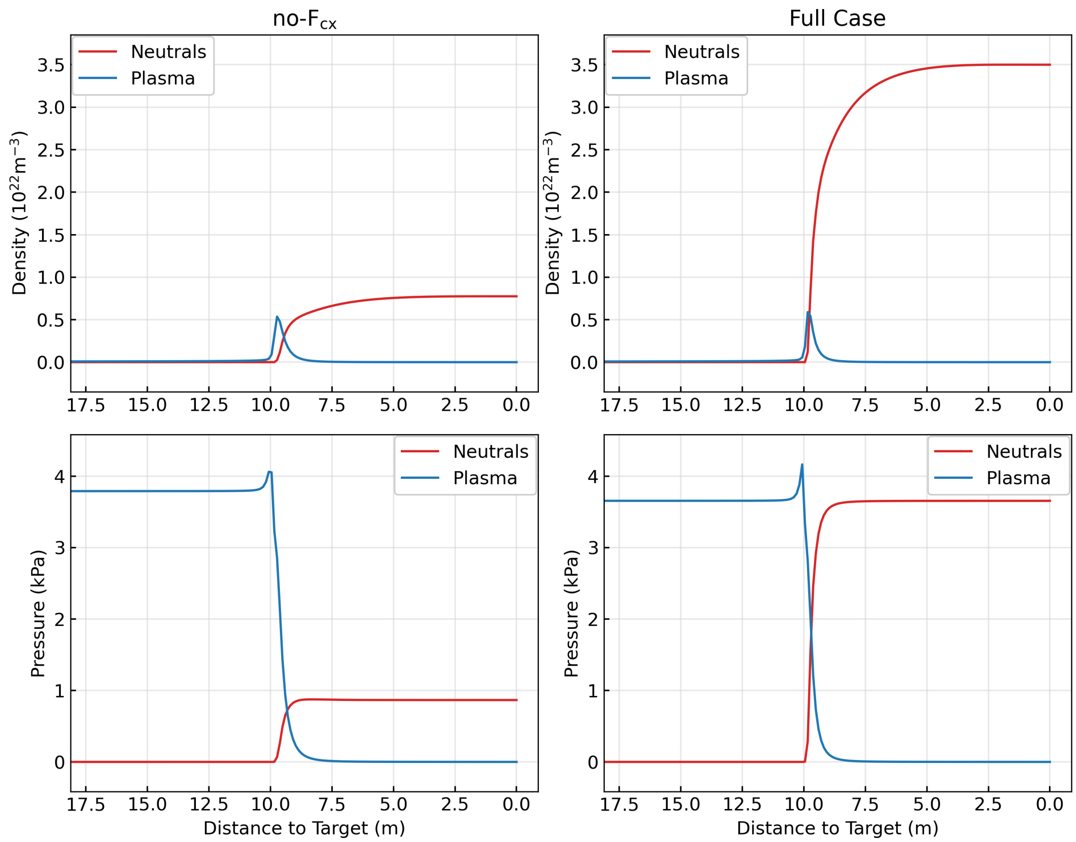

Figure 3. Steady state profiles for the no-F_{CX} (left) and full case (right), depicting the density (top) and pressure (bottom) profiles. The upstream density was set to 4.368 and 4.175 m^{-3} for the no-F_{CX} and full cases, respectively. Neutral-plasma pressure balance is only established in the full case where F_{CX} is included, increasing neutral density. The origin of the peak in plasma pressure is currently uncertain and under investigation.

detachment front position appears to be largely insensitive to the inclusion of F_{CX} where only a small $< 5\%$ decrease in upstream density was required to achieve the same front position.

It can be seen in Figure 3 that for the no-F_{CX} case, the plasma and the neutrals are far from the pressure balance that is seen for the full case. In the no-F_{CX} case, the deficit in the pressures is accounted for by the momentum sink of the charge exchange reactions; the plasma is still losing its momentum and so experiences the same resistive forces, but this does not get passed to the neutrals. Additionally, the peak in the plasma density preceding the front is similar to the density in the neutral gas cloud. For the full case, the neutrals now provide the resistive force themselves and so their pressure increases to match that of the plasma. In order to reach this pressure without increasing temperature, a much higher density is required than before. The gas cloud increases in pressure from $\sim 1\,\mathrm{kPa}$ to $\sim 3.5\,\mathrm{kPa}$ and so density must rise by a similar factor. These pressures and densities are far higher than what is seen in contemporary devices and what is expected in future reactors.

The 1D nature of these simulations, coupled with no particle loss at the target, results in the neutrals being unrestricted in reaching these high pressures and densities. The testing of an imposed neutral particle sink to emulate a neutral gas pump has shown that the neutral pressure can be significantly reduced, but at the cost of requiring radically upstream conditions. Further work is needed in implementing neutral particle sinks, for example through neutral reservoirs [28].

In the present work, the ion target recycling was assumed to be completely thermal, resulting in a complete loss of ion momentum to the target. The neutral momentum boundary condition was set to $V_n = 0$ for both the target and the midplane. In reality, a portion of the ions will be fast recycled into neutrals with non-zero momentum. A more physically correct neutral target boundary should allow some slip since they are not bound to the magnetic field and have a non-negligible momentum in the toroidal direction. Therefore, a more realistic treatment of near-wall momentum could result in the increase the dynamic pressure of neutrals and the target, and reduce static pressure.

4.2. Semi-Empirical Reattachment Model

A semi-empirical model for the response of the detachment front position ($s_{\text{front}}(t)$) to an increase in power was proposed by Henderson et al. [35, 17], given by the following equation 7.

$$s_{\text{front}}(t) = s_{\text{front}}(0) \exp\left(\frac{-tP_{\text{sep}}(t)}{E_{\text{neut}}\frac{\tau_{\text{resid}}}{\tau_{\text{ionis}}}}\right) \tag{7}$$

Here, $s_{\text{front}}(0)$ is the starting front position, E_{neut} is the total energy required to ionise the neutral cloud, given by $E_{\text{neut}} = V_{\text{div}} n_0 E_{\text{ion}}$ where V_{div} is the divertor volume, n_0 the neutral gas density, and $E_{\text{ion}} \sim 30\,\text{eV}$ is the ionisation energy. P_{sep} is used here as the additional power crossing the separatrix over the steady state value, instead of the total power as used in the original definition, so that the model does not predict reattachment for steady-state conditions. The values τ_{resid} and τ_{ionis} are the hydrogenic residence and ionisation timescales, respectively. The form of the relationship being an exponential was chosen empirically to match the functional form of the front movement as observed in experiments. The exponent itself represents the ratio of total input energy to the total amount of energy needed to burn through the neutral cloud. The timescale ratio can be thought of as representing the number of ionisation-recombination cycles a hydrogen atom can undergo during its residence time in the divertor volume [36], acting as a multiplicative factor that raises the total amount of energy needed. An approximation for this ratio is given as the following equation,

$$\frac{\tau_{\text{resid}}}{\tau_{\text{ionis}}} = \frac{L_{||} n_e S_{CD}}{\sqrt{2eT_e/m_D}}. \tag{8}$$

$L_{||}$ is the connection length between the X-point and the target. The model was originally formed for detached divertors where the front had progressed fully to the

X-point, such that to reattach, the front would need to travel the full length of the flux tube within the divertor. For partial detachments studied here, the initial front location was used instead. S_{CD} is the hydrogenic ionisation rate coefficient, evaluated for the values of T_e and n_e at the ionisation front. The ionisation front is defined here as the point of peak ionisation rate, a readily extractable value from the simulation outputs. Different definitions are explored in Section Appendix A. m_D is the mass of a deuterium ion.

The model and its applicability to general detachment front motion is tested across two types of pulses: 1) a rapid step up in power, and 2) a slow rise, comparable to the experiments used as the basis for the original development of the model.

4.3. Rapid Power Step Response

As the semi-empirical model is based entirely on burn-through, we expect it to agree well with no-F_{CX} simulations of transients, since burn-through will dominate the front motion. Figure 4 shows the front response to fast rises in input power. Indeed, the semi-empirical model is able to predict well the form of the front movement as well as the reattachment timescale to within a factor of two. For the 1.4× case, the increase in power is not enough to reattach and so the model cannot predict this behaviour.

Figure 5 depicts the front response to the same power transients for the full case simulations. With the F_{CX} term now enabled, pressure effects now impact the front response and two distinct phases can be seen. The exponential fits from the semi-empirical model are unable to reproduce the exact dynamics found in the simulation because they only follow a single characteristic time. Despite this, the model is still able to largely reproduce the same reattachment time. The disagreement between the specific dynamics of the simulations and semi-empirical model is to be expected, as an instantaneous step up in power is well outside the data set that was used to form the model, as are any pressure balance effects. However, the model's ability to predict the reattachment time might not be, since this is controlled by burn-through alone, which remains unaffected by the pressure effects.

In the first phase of the front response, the movement of the front position is now far quicker, responding directly to the increased power on a scale that appears instantaneous in the time resolution considered here. The front movement then significantly slows down once the power stops rising and reaches the flat-top. The former of these responses is a pressure-dominated movement, while the latter is dominated by burn-through, related to the energy available for ionisation.

An increase in input power is fundamentally an increase in the upstream pressure, i.e. energy density, that acts to raise the plasma pressure within the entire flux tube. With momentum transfer enabled, the plasma applies an increased pressure on the neutrals. The neutral cloud is compressed and the front approaches the target. For rise times of ~ 4 ms and greater, as shown here, the plasma and neutrals are able to respond on the same timescale as this change in power, so the front movement closely matches

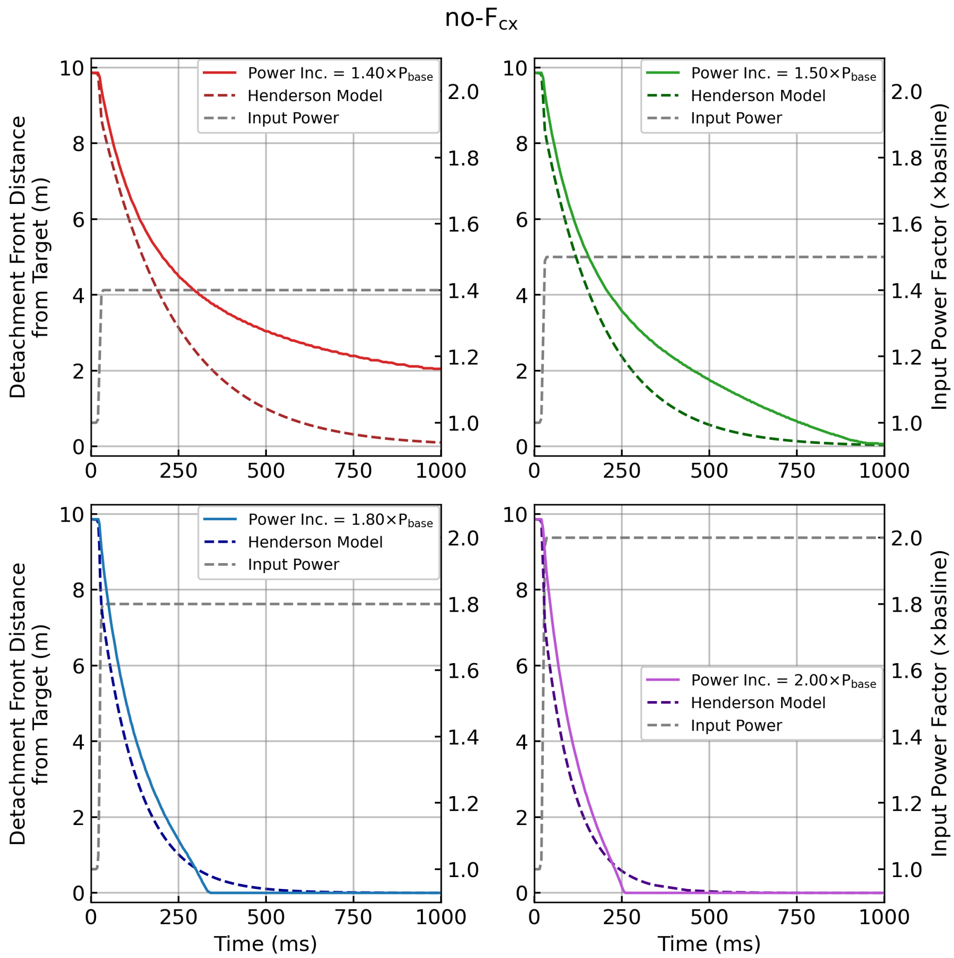

Figure 4. Evolution of the detachment front position in a no-F_{CX} (i.e. burn-through only) simulation in response to a rapid step up in power into the SOL. The power rise (grey dashed line) is shown as a factor of steady-state given in the legends. Hermes-3 outputs are plotted by solid lines, and the appropriate semi-empirical model of Section 4.2 is plotted in dashed lines. The response to a factor of 1.4 finds a new steady state at ∼1.8 m, the others reattach.

the transient's time history. The semi-enclosed nature of the flux tube cylinder results in the compressive force of the plasma acting like a piston on the neutral gas cloud.

At the point of peak power, and thus peak compression of neutrals, the front speed reduces significantly. A steady power input results in a similarly steady plasma pressure and the front movement becomes solely determined by the burn-through of the neutral cloud, i.e. the speed is determined by the net rate of ionisation over recombination. However, notice that the burn-through speeds of the front movement in the full and no-F_{CX} cases differ considerably. This is due to the total number of neutral particles in the full case being significantly higher than the no-F_{CX} case, as shown in Figure 1. For equivalent ionisation rates, more neutral atoms must be ionised for the full case to produce the same front displacement. Nevertheless, even with this extended timescale, the rapid step to 2× power in the full case still reattaches at ∼1.2 s, approximately at

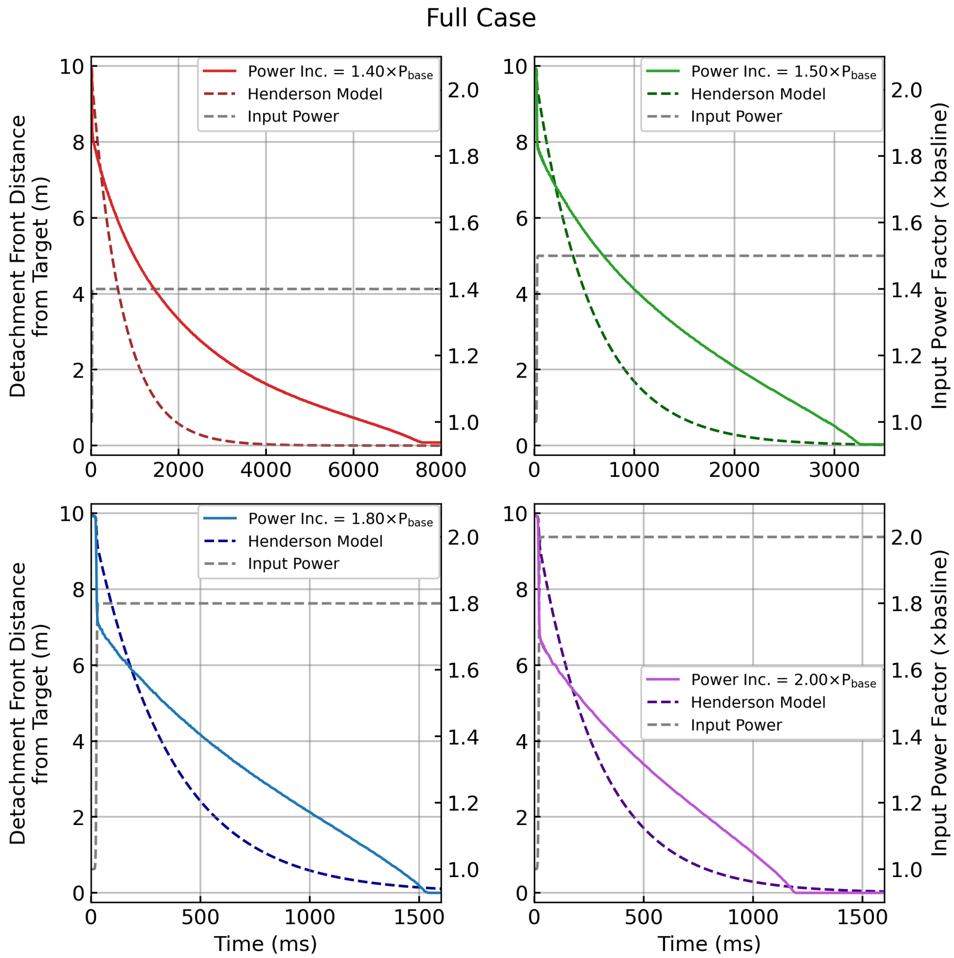

Figure 5. Evolution of the detachment front position in the full case (i.e. including F_{CX}) in response to a rapid step up in power into the SOL. The power rise (grey dashed line) is shown as a factor of steady-state given in the legends. Hermes-3 outputs are plotted by solid lines, and the appropriate semi-empirical model of Section 4.2 is plotted in dashed lines. The 1.5× case reattaches after ∼3 s, and the 1.4× case requires >7 s. The semi-empirical model has a smaller exponent for the full case than the no-F_{CX} case due to the increased initial number of neutrals.

the same time as the no-F_{CX} case at a factor of 1.5× baseline.

4.4. Slow Ramp Response

While the rapid steps in power are useful for understanding pulse response, the semi-empirical model was developed by comparing to experimental data with a much slower change in power. No-F_{CX} simulations featuring similar power transients are shown in Figure 6, with the power increasing over 200 ms to a given factor of the baseline.

There is reasonable agreement between the simulations and model for cases that reattach, but predictably not for the 1.2× case, which does not. The form of the model matches closely with the simulations, though it is apparent that the exponent is slightly larger for the model, over-predicting the rate of burn-through. Despite the model not

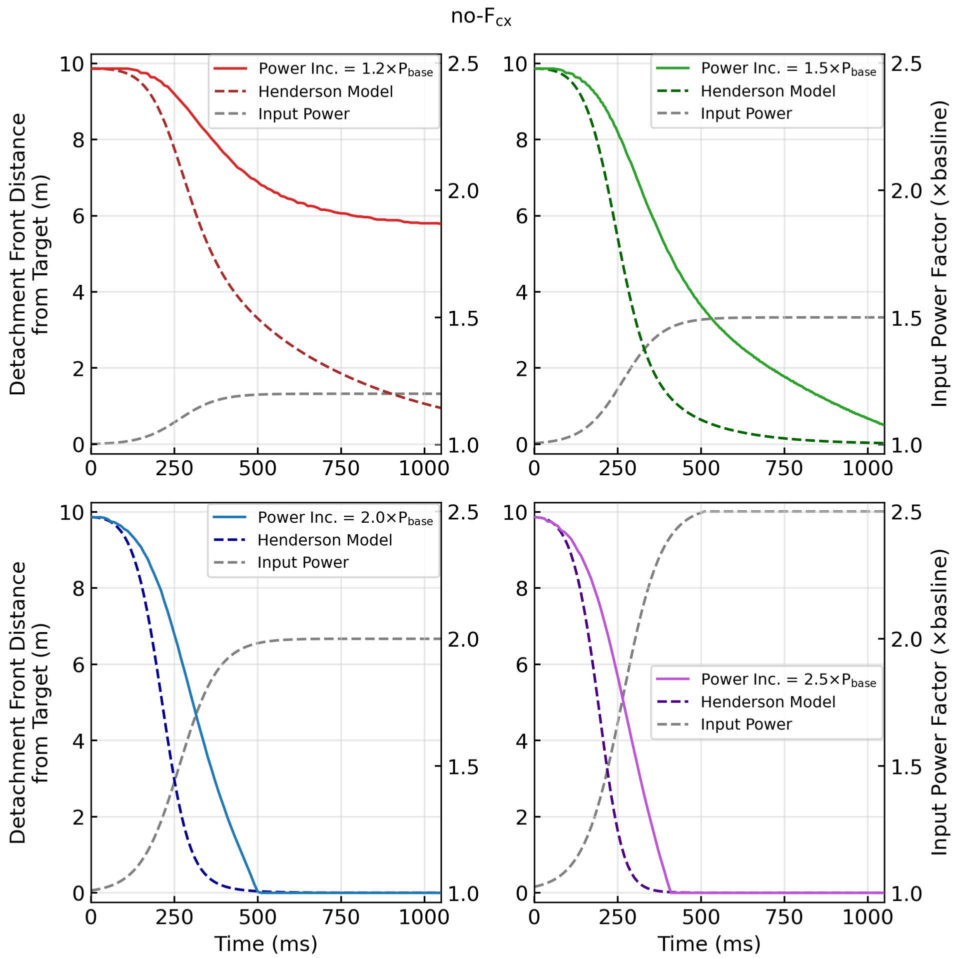

Figure 6. Evolution of the detachment front position in a no-F_{CX} simulation in response to a slow ramp up in power into the SOL. The power ramp is plotted as a factor of steady-state (with axis on the right). Hermes-3 outputs are plotted by solid lines, and the appropriate semi-empirical model of Section 4.2 is plotted in dashed lines. The 1.2× case finds a new steady state of ∼5 m, while all others reattach.

predicting a finite reattachment time, there is also reasonable agreement for when the front reaches very close to the target, i.e. when it is close to reattaching, improving for higher powers.

In contrast to the rapid step cases in Section 4.3, the full case simulations that are shown in Figure 7 no longer show clear evidence of two distinct mechanisms of front movement. The front movement is still being impacted by the changing pressure, but the rise in power is slow enough that the magnitude of this effect is comparable to that of the burn-through, and it is less distinguishable. A subtle but clear change in the front speed can be seen as the rise in power begins to level off, highlighting that the pressure-driven compression of the neutral cloud is still present. The choice of powers that were explored are the result of this effect, as the powers shown here provides scenarios where the plasma reattaches ranges from timescales larger than the power increase timescale to those where they are approximately equal, as is seen for the no-F_{CX} cases. These

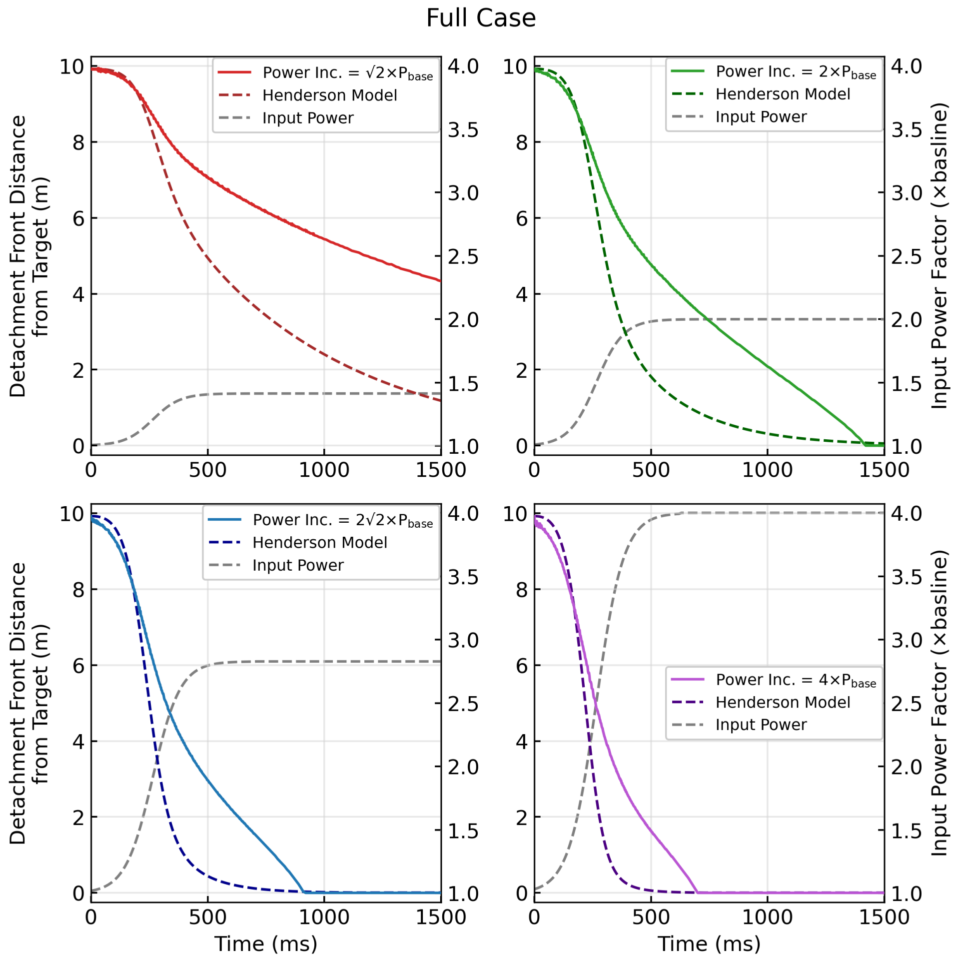

Figure 7. Evolution of the detachment front position for the full case in response to a slow ramp up in power. The power ramp is plotted as a factor of steady-state (with axis on the right). Hermes-3 outputs are plotted by solid lines, and the appropriate semi-empirical model of Section 4.2 is plotted in dashed lines. Note the subtle change in slope at the end of the pulse where only burn-through is in effect. Front movements approximately follow the power ramp, making experimental differentiation between pressure- or ionisation-dominated regimes difficult.

higher powers are outside the moderate $\sim 2\times$ powers upon which the semi-empirical model is based, yet the agreement is still within a reasonable margin of error.

Similar to the no-F_{CX} cases, the simulations shown in Figure 7 compare well with the model, both in terms of reattachment times and dynamic form, largely for the earlier part of the transient though. Again, the model over-predicts the exponent value, suggesting faster reattachments. The greater neutral density results in longer reattachments due to the increased ionisation energies required.

Similar steady state positions are reached by the full and F_{CX} cases for equivalent low power rises that do not reattach, reaffirming that the F_{CX} term has only a minor impact on the steady state front position.

The forms of the full and no-F_{CX} cases are no longer as distinct from one another as they are for the rapid power rises. This is of particular relevance when comparing with the semi-empirical model, as power rises of the forms given in Figures 6 & 7 are similar to those in the experimental studies on which the model is based. As the power increases much more slowly, the plasma-neutral pressure balance is still reached quasi-instantaneously. However, now the timescale of the burn-through is comparable to the timescale of the power rise, both in the few 100 ms range, and similar to typical experimental energy confinement times, τ_{E}. In such cases, the magnitudes of burn-through and neutral compression become similar and indistinguishable from each other, both in simulation and experiment, which provides a possible explanation as to why the model does not need a specific provision for pressure effects.

The match between the Henderson model and the Hermes-3 simulations needs to be qualified by considering the variation in the value of S_{CD} with temperature. The outline of the model suggests a range of $2-3\,\mathrm{eV}$ to be appropriate, where T_e and n_e can effectively be used as tuning parameters. The values found from the simulations presented here, however, are outside that at $\sim 1.6\,\mathrm{eV}$. Choosing the relevant plasma parameters to provide the closest fit could be achieved through using different definitions of the ionisation front. Using instead the point of the detachment front gives significantly lower exponent values by an order of magnitude. Refinements to the model might additionally seek to constrain the timescale ratio in the denominator more accurately.

4.5. Finite Length Pulses

Finally, the effect of finite-length transient pulses on the detachment front position has been investigated.

Figure 8 shows the detachment front position response in both types of simulations to a pulse with a rise time of 4 ms and a duration of 50 ms, along with the net plasma ionisation rate, i.e. ionisation minus recombination.

The initial difference in front position at $t = 0$ between full and no-F_{CX} cases is the result of an imperfect rebalance of the front position due to the inclusion of F_{CX}, and has a small enough impact on the front dynamics to be ignored. The front movement during the power rise is very similar to that shown in Figures 4 and 5 for the fast power rises. For the full case simulations, the front tracks the power rise instantaneously at first and is followed by a burn-through that is slower than that in the no-F_{CX} case. However, it can be seen from the right-hand plot of Figure 8 that the net ionisation rate during the rise and peak of the power pulse is very similar between the two cases, meaning that ionisation and burn-through are occurring throughout the entire pulse at a rate that is largely insensitive to the simulation type and the instantaneous front position.

Once the power returns to the baseline value, the net ionisation rate drops below zero and net recombination occurs. For the full case, the high-pressure neutral gas pushes back on the plasma and the front retreats. Since a finite number of neutrals have been ionised, the total compressive energy stored in the gas cloud is lower, and so

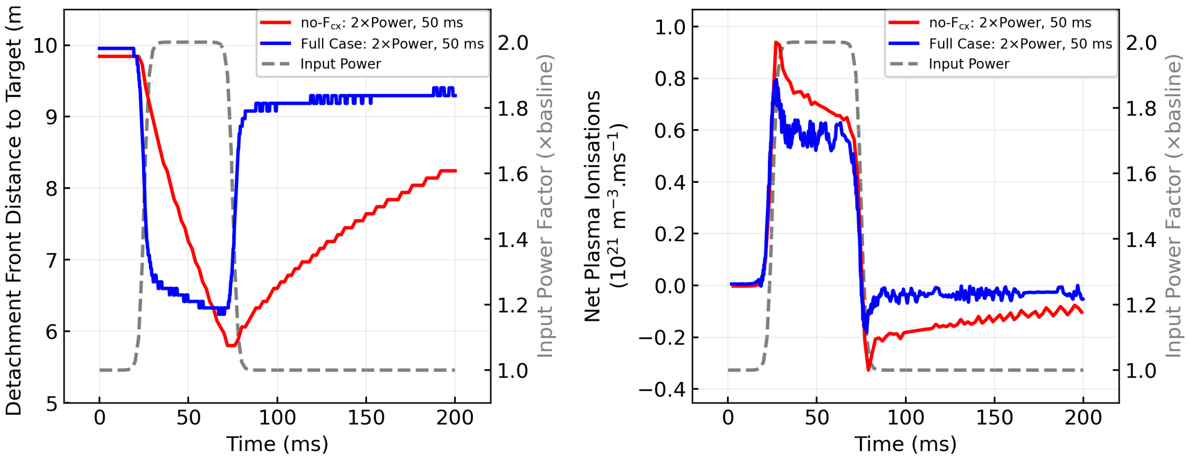

Figure 8. (left) The response of the detachment front position due to a 50 ms transient for both full and no-F_{CX} cases. (right) The rate of net ionisations during this pulse. Despite the initial much faster front movement for the full case compared with the no-F_{CX}, they both have very similar ionisation rates.

the front does not return fully to its original position. For both cases, recombination steadily refills the gas cloud, and the front will return to its original position on a timescale of many 100's ms. The rate at which the neutrals recombine is significantly reduced compared with the ionisation rate during the increased power input, resulting in the slower front movement for both cases. The slower front movement of the full case, resulting from recombination after the fast neutral expansion, is in part expected due to the previously mentioned higher neutral density requiring a larger amount of recombination for the same front speed. However, the lower recombination rate is currently unexplained. A potential justification could be that the neutral expansion returns the front close to its steady state position much faster than the comparatively slow rate for the no-F_{CX} case. The smaller recombination rate then occurs because of the front is far closer to its ultimate steady state position.

For both cases, the front position does not return to its initial steady state position within the energy confinement of ∼100 ms, meaning that repeated increases in power of this form would result in a progressive movement of the front towards the target. The rapid retreat of the front in the full case results in a smaller initial hysteretic effect than the no-F_{CX} case.

Finally, to connect to the description of both the semi-empirical model and the steady-state simulations above, it is worth noting that the net ionisation rate is $\sim 1\%$ of the total ionisation rate, indicating that the front movement is the result of a very small imbalance on the steady-state conditions.

5. Discussion & Conclusions

This study has investigated the dynamic response of detached plasmas to power transients in a STEP-like reactor configuration utilising 1D Hermes-3 fluid simulations.

By examining cases with and without neutral momentum transfer via charge exchange, we have gained valuable insights into the physics governing detachment stability during transient events.

Our results demonstrate that detachment front dynamics are governed by two distinct mechanisms: neutral burn-through and plasma-neutral pressure balance. In the absence of neutral momentum transfer, no-F_{CX}, front movement is dominated solely by the burn-through process, resulting in a relatively slow response that approximately follows an exponential temporal evolution. For fast power rises with neutral momentum included, a two-phase response emerges: an initial rapid compression of the neutral gas cloud driven by plasma pressure, followed by a slower burn-through phase.

The semi-empirical reattachment model proposed by Henderson et al. performs remarkably well for slow power ramps. It is important to note that due to the close match of the front to the slowly increasing pressure in momentum-enabled cases, it may not be possible to differentiate between burn-through and pressure-dominated regimes in the experimental data, such as those that the semi-empirical model is based on. For rapid power steps, significant deviations in front movement form from model predictions were observed but similar reattachment times were predicted. This highlights how a pure of a purely burn-through based approach may provide utility. The model's performance deteriorates further in the full cases when F_{CX} is included, as it does not capture the rapid compression of the neutral gas cloud, thus underscoring the need for more comprehensive analytical frameworks that account for pressure-driven dynamics.

Our investigation of finite-length pulses revealed hysteretic behaviour in detachment front position for transients with timescales below $\sim 100\,\mathrm{ms}$, with full case simulations showing reduced hysteresis due to the neutral expansion compared with the no-F_{CX} case. This finding has significant implications for operational scenarios involving repeated transients, such as vertical displacement events, since it could result in a progression of the average front position with each repeated event.

It is acknowledged that neglecting charge exchange momentum transfer to the neutrals seems unphysical, and yet much of the dynamics seen in the semi-empirical model, and thus the experimental data on which it is based, is best approached in this scenario. However, note that the 1D nature of the simulations presented here implies significant restrictions on the realism of the output. In particular, the true nature of the neutral gas dynamics is kinetic, and thus a fluid approximation will have important shortcomings. To be precise, full momentum transfer is expected to be a poor approximation because the neutrals are not confined to the magnetic field line as a 1D flux tube would imply. This would result in transverse motion of the gas under compression, meaning that not all of the neutrals will contribute to the pressure-driven effects. On the other hand, the physical interpretation of the no-F_{CX} case is that neutrals, which take on the momentum of rapid plasma particles, immediately leave the domain. The important insight here, therefore, is that a realistic dynamic solution is expected to lie between these two limiting cases, on which our simulations have now put some clear bounds.

The substantial difference in neutral gas density between the simulation cases emphasises the importance of accurate neutral modelling in detachment studies. The pressure balance established in the full cases leads to much higher neutral densities than typically expected in tokamak divertors, suggesting that additional mechanisms for neutral pressure relief may play an important role in real devices, and the need for more sophisticated modelling in this area.

The 1D approach used in this study provides valuable insights into detachment dynamics while enabling high-throughput simulations suitable for parameter scans. However, limitations in the treatment of cross-field transport and the closed particle system necessitated the use of a PI controller to emulate particle losses. Future work should focus on extending these simulations to include physics-based neutral particle sinks such as reservoir models and more realistic 2D effects, particularly the radial spreading and drift of the peak heat flux channel, which could significantly affect the effective flux expansion. Full 2D simulations time-dependent simulations are planned to investigate this effect.

These findings contribute to our understanding of detachment control in future fusion devices like STEP. The identified two-timescale response to power transients suggests that detachment control systems may need to address both rapid pressure-driven effects and slower burn-through dynamics. Additionally, the observed hysteresis in detachment front position for short pulses highlights the importance of considering the history of transient events when developing control strategies for maintaining detachment during high-performance operation.

Acknowledgments

Prepared in part by LLNL under Contract DE-AC52-07NA27344

This work has been funded by STEP, a UKAEA program to design and build a prototype fusion energy plant and a path to commercial fusion.

This project was funded by EPSRC CDT in the Science and Technology of Fusion Energy, Grants EP/L01663X/1, EP/S022430/1.

Appendix A. Model Exponent Sensitivity

The value of the parameters needed for the model exponent are taken from the simulation outputs. The hydrogenic ionisation rate coefficient, S_{CD}, is calculated using the values of n_e and T_e found at the ionisation front in the initial steady state conditions. In the work given above, the ionisation front is defined as the point of peak ionisation rate, as this is how the model itself was designed. There is, however, a strong sensitivity in the overall exponent value to the location where these parameters are extracted.

Figure A1 depicts four cases, split between full/no-F_{CX} and fast/slow power rises, of front response to power changes and the model predictions for various ionisation front definitions. The ionisation peak is typically found at a point where the temperature

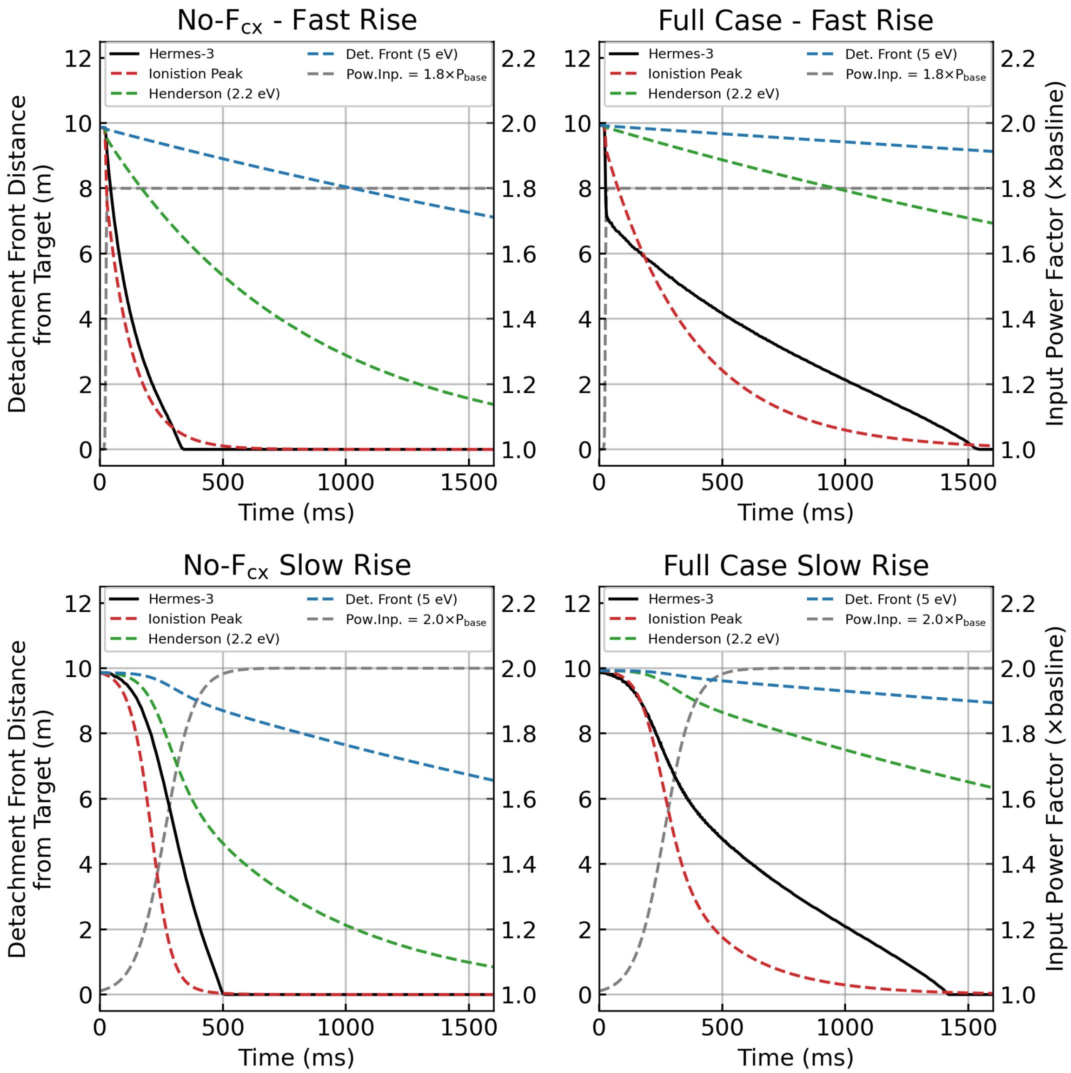

Figure A1. Detachment front response to power rises, and predictions from the semi-empirical model for different front definitions. (left) No-F_{CX} cases, (right) full cases, (top) 4 ms fast power rise, (bottom) 200 ms slow power rise. The ionisation peak is the point of peak ionisation rate, Henderson is evaluated at the point where the temperature is 2.2 eV, and the detachment front is defined as the point of 5 eV.

is ∼1.6 eV, lower than the typical range of 2-3 eV suggested during the model's development. A value of 2.2 eV was given as the expected value in the model definition, whereas the 5 eV point is the point of the detachment front. In all four cases, the ionisation peak provides the best agreement with the simulations. Though it appears that an exponent value that would give the closest alignment would be somewhere between the ionisation peak and the 2.2 eV value. The spread of potential fits highlights that the choice of temperature can act as a tuning parameter for the model, as Henderson et al. originally suggested. Given the uncertainties associated with the simulations themselves, and with the model's formation, the remarkable agreement seen here should be very encouraging, though not a definitive demonstration of the accuracy of either.

Appendix B. Influence of the Upstream Density Controller

In the present work, the PI controller is used as a way to include cross-field transport effects. Since this is only an emulation of physical transport, its impact is assessed through additional simulations.

Figure B1 shows the detachment front response of a no-F_{CX} case to a power pulse for different controller settings. The power pulse has a width of 50 ms, rise and fall times of 4 ms, and peak power of 2×baseline. Shown also is the upstream density to highlight the role of the PI controller. They all have the same specified target upstream density and the ratio between the proportional and integral terms is kept constant.

With the PI controller off, the upstream density significantly increases as the power rises and remains high throughout the pulse. There is no removal of plasma particles and thus the density increases, raising the hydrogenic and impurity radiation that resists reattachment. This results in a reduced ionisation rate and slower front progression. This extreme case of the PI controller being fully off produces unphysically resilient detachments, as discussed in Section 3.2. The speed of the front to return to the steady state position is also far slower compared with the other controller settings, likely due to the lack of a particle source.

For the low and high settings, the upstream density still rises with the power increase, but does not remain high. The controller is removing particles to bring the density down to the set level, but the number of plasma particles created from ionisation is rising faster than the weaker gain terms can respond to. Once the power returns to baseline, the upstream density falls below the set value, and the controller acts as a particle source. These intermediate controller settings appear to have delayed responses to the power changes compared with the maximum and off cases. This can be seen in the overshoot and corrections in the upstream densities for the low and high settings after the pulse has finished, hence why the low setting continues to move inwards even after the power has returned to baseline.

For the maximum controller setting, the upstream density remains at the set level throughout. This provides the maximum particle removal rate, so the front speed and displacement are the highest. The front also responds immediately to the fall in power, displaying no oscillatory behaviour.

Figure B2 depicts the front position and upstream density against time for the full case. The two-phase behaviour is still seen for all controller settings, with the magnitude of the initial rapid movement the same for all settings, reinforcing that this is a pressure-driven phenomenon only. The rate of front progression during the power pulse flat-top increases with controller settings, as expected, as this is now a burn-through dominated regime. The magnitude of the pressure-driven return is also dependent on the controller setting, with higher removal rates producing a weaker return as more particles have been removed and thus less compressive energy is stored in the neutrals. Similar oscillatory behaviour for the intermediate controller settings is seen here as in the no-F_{CX} cases.

If given enough time, any non-zero PI controller settings will result in the upstream

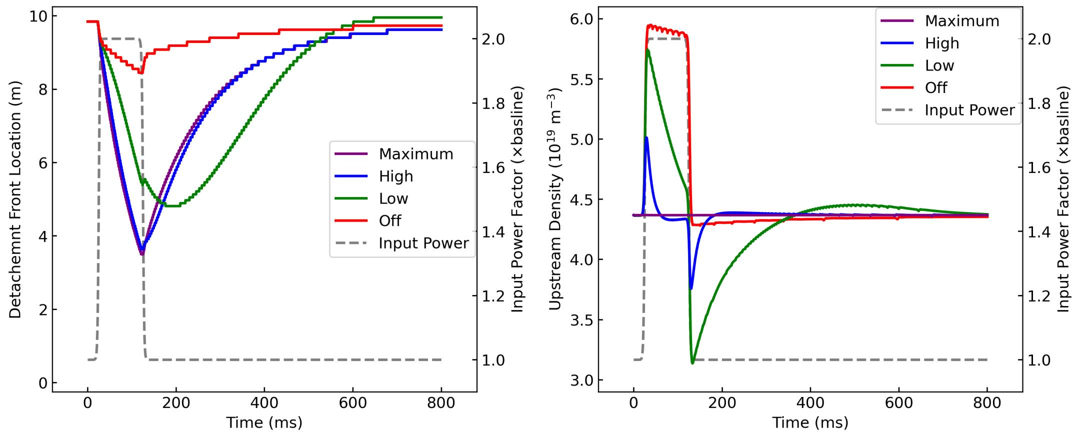

Figure B1. No-F_{CX} simulation results, (left) detachment front position against time as a response to a power pulse for different PI controller settings, (right) upstream density against time.

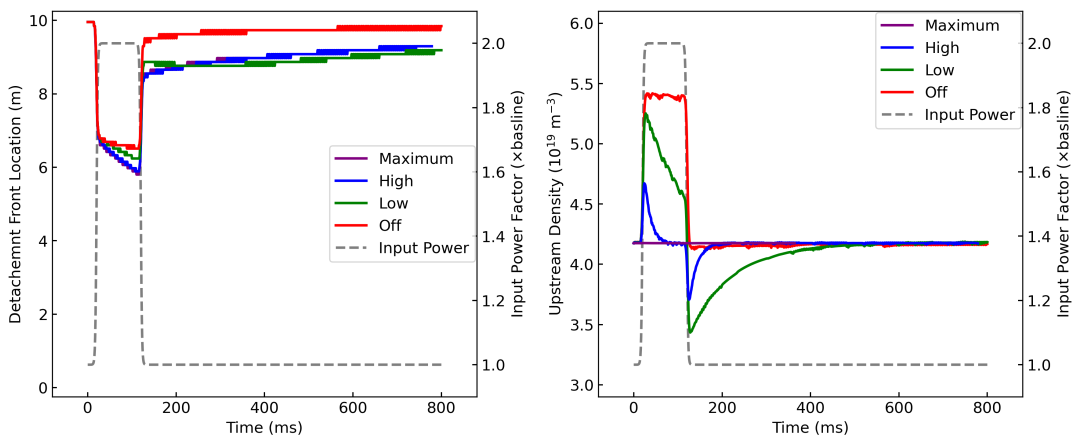

Figure B2. Full case simulation results, (left) detachment front position against time as a response to a power pulse for different PI controller settings, (right) upstream density against time.

density reaching the set value, with weaker gain settings resulting in higher initial peaks and slower times to reach the set level. A priori, it is unclear whether the upstream density should rise at all, and if so, by how much. This leaves the gain terms as tuning parameters that could be used to more closely match higher-fidelity simulations or experiments. For fast rise times, the same qualitative comparison is seen. The fast front movement is still seen, but without the particle removal, the burn-through progression of the front is far slower.

[1] Wilson, Howard and Chapman, Ian and Denton, Tris and Morris, William and Patel, Bhavin and Voss, Garry and Waldon, Chris and the STEP Team, 2020 STEP-on the pathway to fusion commercialization *Commercialising Fusion Energy* 2053-2563 (IOP Publishing) pp 8–1 to 8–18 ISBN 978-0-7503-2719-0 URL `https://dx.doi.org/10.1088/978-0-7503-2719-0ch8`

[2] Muldrew S I, Harrington C, Keep J, Waldon C, Ashe C, Chapman R, Griesel C, Pearce A J, Casson F, Marsden S P *et al.* 2024 *Fusion Engineering and Design* **201** 114238

[3] Hudoba A, Newton S, Voss G, Cunningham G and Henderson S 2023 *Nuclear Materials and Energy* **35** 101410

[4] Tholerus E, Casson F, Marsden S, Wilson T, Brunetti D, Fox P, Freethy S, Hender T, Henderson S, Hudoba A *et al.* 2024 *arXiv preprint arXiv:2403.09460*

[5] Mario Merola and D Loesser and A Martin and P Chappuis and R Mitteau and V Komarov and RA Pitts and S Gicquel and V Barabash and L Giancarli and J Palmer and M Nakahira and A Loarte and D Campbell and R Eaton and A Kukushkin and M Sugihara and F Zhang and CS Kim and R Raffray and L Ferrand and D Yao and S Sadakov and A Furmanek and V Rozov and T Hirai and F Escourbiac and T Jokinen and B Calcagno and S Mori 2010 *Fusion Engineering and Design* **85** 2312–2322 ISSN 0920-3796 proceedings of the Ninth International Symposium on Fusion Nuclear Technology URL `https://www.sciencedirect.com/science/article/pii/S0920379610004060`

[6] Jackson G, Politzer P, Humphreys D, Casper T, Hyatt A, Leuer J, Lohr J, Luce T, Van Zeeland M and Yu J 2010 *Physics of Plasmas* **17**

[7] Cohen R H and Xu X 1995 *Physics of Plasmas* **2** 3374–3383

[8] Baylor L R, Commaux N, Jernigan T C, Brooks N, Combs S K, Evans T, Fenstermacher M, Isler R, Lasnier C, Meitner S *et al.* 2013 *Physical review letters* **110** 245001

[9] Kool B, Verhaegh K, Derks G, Wijkamp T, Lonigro N, Doyle R, McArdle G, Vincent C, Lovell J, Federici F *et al.* 2024 *arXiv preprint arXiv:2407.07784*

[10] Osawa, RT and Moulton, D and Newton, SL and Henderson, SS and Lipschultz, B and Hudoba, A 2023 *Nuclear Fusion* **63** 076032

[11] Leonard A W 2014 *Physics of Plasmas* **21**

[12] Wiesen S, Reiter D, Kotov V, Baelmans M, Dekeyser W, Kukushkin A, Lisgo S, Pitts R, Rozhansky V, Saibene G, Veselova I and Voskoboynikov S 2015 *Journal of Nuclear Materials* **463** 480–484 ISSN 0022-3115 pLASMA-SURFACE INTERACTIONS 21 URL `https://www.sciencedirect.com/science/article/pii/S0022311514006965`

[13] Park J S, Lore J D, Reinke M L, Kuang A Q, De Pascuale S and Creely A J 2024 *Nuclear Fusion*

[14] Lipschultz B, Parra F I and Hutchinson I H 2016 *Nuclear Fusion* **56** 056007

[15] Cowley C, Lipschultz B, Moulton D and Dudson B 2022 *Nuclear Fusion* **62** 086046

[16] Myatra O, Lipschultz B, Moulton D, Verhaegh K, Dudson B, Orchard S, Fil A and Cowley C 2023 *Nuclear Fusion* **63** 096018

[17] Henderson S, Bernert M, Brida D, Cavedon M, David P, Dux R, Février O, Jacquet P, Järvinen A, Kallenbach A *et al.* 2024 *Nuclear Fusion* **64** 066006

[18] Dudson B, Kryjak M, Muhammed H, Hill P and Omotani J 2024 *Computer Physics Communications* **296** 108991

[19] Body T, Eich T, Kuang A, Looby T, Kryjak M, Dudson B and Reinke M 2024 *Nuclear Materials and Energy* **41** 101819

[20] Dudson, B D and Allen, A and Breyiannis, G and Brugger, E and Buchanan, J and Easy, L and Farley, S and Joseph, I and Kim, M and McGann, A D and et al 2015 *Journal of Plasma Physics* **81** 365810104

[21] Hindmarsh, Alan C and Brown, Peter N and Grant, Keith E and Lee, Steven L and Serban, Radu and Shumaker, Dan E and Woodward, Carol S 2005 *ACM Trans. Math. Softw.* **31** 363–396 ISSN 0098-3500 URL `https://doi.org/10.1145/1089014.1089020`

[22] Reiter D *et al.* 2000 *Forschungszentrum Juelich GmbH* **52425**

[23] Kotov V, Reiter D and Kukushkin A S 2007 Numerical study of the ITER divertor plasma

with the B2-EIRENE code package Tech. rep. Forschungszentrum Juelich (Germany). Inst. fuer Energieforschung (IEF)

[24] Moulton D, Stangeby P, Bonnin X and Pitts R 2021 *Nuclear Fusion* **61** 046029 URL `https://dx.doi.org/10.1088/1741-4326/abe4b2`

[25] Stangeby P 1995 *Physics of plasmas* **2** 702–706

[26] Dudson B D, Allen J, Body T, Chapman B, Lau C, Townley L, Moulton D, Harrison J and Lipschultz B 2019 *Plasma Physics and Controlled Fusion* **61** 065008

[27] Henderson S, Osawa R, Newton S, Moulton D, Xiang L, Futtersack R, Kryjak M, Ridgers C, Karhunen J, Jarvinen A, Hudoba A, Bakes S, Eriksson F, Meyer H, Lord M, Tarazona A, Cureton A, Barth A, Chuilon B, Hebrard T, Wang S, Vizvary Z, Vaccaro D, Perez Smith F, Farrington J, Harrison J, Dudson B and Lipschultz B 2024 *Nuclear Fusion* **65** 016033 URL `https://dx.doi.org/10.1088/1741-4326/ad93e7`

[28] Derks G, Frankemölle J, Koenders J, van Berkel M, Reimerdes H, Wensing M and Westerhof E 2022 *Plasma Physics and Controlled Fusion* **64** 125013

[29] Cowley C, Moulton D and Lipschultz B 2024 *arXiv preprint arXiv:2407.13501*

[30] A Loarte and B Lipschultz and AS Kukushkin and GF Matthews and PC Stangeby and N Asakura and GF Counsell and G Federici and A Kallenbach and K Krieger and A Mahdavi and V Philipps and D Reiter and J Roth and J Strachan and D Whyte and R Doerner and T Eich and W Fundamenski and A Herrmann and M Fenstermacher and P Ghendrih and M Groth and A Kirschner and S Konoshima and B LaBombard and P Lang and AW Leonard and P Monier-Garbet and R Neu and H Pacher and B Pegourie and RA Pitts and S Takamura and J Terry and E Tsitrone and the ITPA Scrape-off Layer and Divertor Physics Topical Group 2007 *Nuclear Fusion* **47** S203 URL `https://dx.doi.org/10.1088/0029-5515/47/6/S04`

[31] Omotani J T and Dudson B D 2013 *Plasma Physics and Controlled Fusion* **55** 055009 URL `https://dx.doi.org/10.1088/0741-3335/55/5/055009`

[32] Wigram M, Ridgers C, Dudson B, Brodrick J and Omotani J 2020 *Nuclear Fusion* **60** 076008 URL `https://dx.doi.org/10.1088/1741-4326/ab868b`

[33] Osawa R, Moulton D, Newton S, Henderson S, Lipschultz B and Hudoba A 2023 *Nuclear Fusion* **63** 076032 URL `https://dx.doi.org/10.1088/1741-4326/acd863`

[34] Nakazawa S, Nakajima N, Okamoto M and Ohyabu N 2000 *Plasma physics and controlled fusion* **42** 401

[35] Henderson S, Bernert M, Brida D, Cavedon M, David P, Dux R, Février O, Järvinen A, Kallenbach A, Komm M, McDermott R, O'Mullane M, the ASDEX Upgrade Team and the EUROfusion MST1 Team 2023 *Nuclear Fusion* **63** 086024 URL `https://dx.doi.org/10.1088/1741-4326/ace2d6`

[36] Janev R and Reiter D 2018 Isotope Effects in Molecule Assisted Recombination and Dissociation in Divertor Plasmas Tech. rep. Plasmaphysik