Statics and dynamics of detachment in 1D

Frankemölle, J.P.K.W.

Award date:
2021

Statics and dynamics of detachment in 1D

Joint Master's graduation project for Applied Physics and Science and Technology of Nuclear Fusion

Jens Peter Frankemölle BSc

Eindhoven, January 2021

Dr. Egbert Westerhof (DIFFER)
Dr. ir. Jan van Dijk (TU/e EPG)
Dr. Hugo de Blank (TU/e FUSION)

“In ancient days two aviators procured to themselves wings. Daedalus flew safely through the middle air across the sea, and was duly honored on his landing. Young Icarus soared upwards towards the sun till the wax melted which bound his wings, and his flight ended in fiasco. In weighing their achievements perhaps there is something to be said for Icarus. The classic authorities tell us that he was only " doing a stunt," but I prefer to think of him as the man who certainly brought to light a constructional defect in the flying machines of his day. So too in science. Cautious Daedalus will apply his theories where he feels most confident they will safely go; but by his excess of caution their hidden weaknesses can not be brought to light. Icarus will strain his theories to the breaking-point till the weak joints gape. For a spectacular stunt? Perhaps partly; he is often very human. But if he is not yet destined to reach the sun and solve for all time the riddle of its constitution, yet he may hope to learn from his journey some hints to build a better machine.”

– Arthur Eddington, 1920

Abstract

The heat flux through the scrape-off layer of a fusion device is so big that it cannot be allowed to pass to the divertor target unmitigatedly. However, mitigation techniques such as target tilting and impurity seeding alone are insufficient, both in the face of ELMs and growing machine sizes. To tackle this, future machines will likely be run with detached divertor plasmas: regimes of greatly reduced target heat and particle fluxes. Dynamic 1D codes, like the DIV1D code that is used in this work, provide an interesting class of models that allow for faster calculations than their bulkier multidynamensional counterparts like SOLPS-ITER, and could prove vital in researching detachment. This thesis looks at two aspects of the DIV1D code in relation to detachment, namely (I) how well it compares to SOLPS-ITER simulations for comparable physics scenarios and (II) how it acts in the face of ELMs perturbing its power balance. Both necessetated the development of novel features in the code: (I) radial power losses and (II) the possibility to add ELMs. These features are described in this work, and their effects on the DIV1D code—particularly in detached regimes—are characterised.

Contents

1. Introduction

1.1 Nuclear fusion and the heat exhaust issue

In its physical essence, nuclear fusion is the phenomenon that allows two atomic nuclei to combine into a third. It is often mentioned in relation to the combined efforts in physics and engineering to build a power plant based on that same principle. Before entering into a discussion of how one goes about building such a device and why the heat exhaust is a particularly nefarious issue in the grand scheme of things, it is advisable to determine why one should want one in the first place. Since fusion electricity demonstration is currently envisaged for early in the second half of this century (EUROfusion, 2018), any proposed reason should at the very least remain relevant up to a minute past midnight on January 1st of 2050. As a result, it does not bode well for at least one promise of fusion. After all, that it is a "source of safe, non-carbon emitting and virtually limitless energy" (ITER Organization, 2020) will effectively be the low bar for any power plant past 2050.

That same date is mentioned in a wholly different setting that is nonetheless intricately linked to nuclear fusion's future: that of global warming. Global warming refers to the experimental observation that temperatures have significantly risen with respect to pre-industrial times (IPCC, 2014). In a summary for policy makers at the end of their fifth assessment cycle synthesis report in 2014 the IPCC note, with respect to its probable cause, that:

> "Their [anthropogenic greenhouse gas emissions'] effects, together with those of other anthropogenic drivers, have been detected throughout the climate system and are *extremely likely* to have been the dominant cause of the observed warming since the mid-20th century"
>
> — IPCC, 2014

In a more recently published special report in parallel to their sixth assessment cycle, the IPCC additionally calculated that to pursue a temperature rise of at most $1.5°C$ requires zero net emissions by 2050 (IPCC, 2018). The Klimaatakkoord (Kabinet-Rutte III, 2019) testifies to the fact that The Netherlands has taken the IPPC's call seriously. It likewise envisions zero emission by power plants by 2050. So zero emissions are the very least that nuclear fusion proponents should bring to the table. Yes, fusion checks the boxes of sustainability—renewable fuel and low emissions—which makes it a viable option for power production in the post-fossil era, unlike coal and gas plants. So do all other renewable energies, one might interject. Nuclear fusion, though, has some considerable advantages over competing technologies that, although not reasons to switch to 100% fusion-based energy *per se*, might allow it a seat at the table. It is unfazed by the whims of nature, unlike solar photovoltaics and wind power (Saarinen, Dahlbäck & Lundin, 2015) and can thus provide a baseload to the grid where one is needed. Furthermore, its high energy density might make it a good candidate in places where land is scarce and there is no place for extensive solar or wind farms which have considerably lower power densities (Van Zalk & Behrens, 2018). Unlike an important other baseload option, nuclear fission, it also does not produce long-lived radioactive waste (Rocco & Zucchetti, 2000). Competing technologies, then, might be phrasing it wrongly. Nuclear fusion, in the energy mix of the future, could be an interesting *partner technology* in situations where other renewables are not at their best.

Unlike for many other power technologies, the seminal formulation of the underlying physics principles is only a century old (exactly so, in October 2020). It was in October of 1920 that Eddington published his "The Internal Constitution of the Stars" in which he set forth a new perspective on stellar power

production via fusion, triggered by recent experiments in subatomic physics and the formulation of a relation between energy and mass deficit (Einstein, 1905). Conclusive evidence for fusion was found by Bethe (Bethe & Bacher, 1936; Bethe, 1937; Livingston & Bethe, 1937). Research started in the UK in 1945, the USA in 1946 and the USSR in 1950 (Shafranov, 2001). A mere quarter century after its earliest conception, fusion had been established as a research field.

Fast-forward 70 years and fusion has yet to leave that experimental stage. There is no one reason for this, as exemplified by the eight missions put forward in the "European Research Roadmap to the Realisation of Fusion Energy" (EUROfusion, 2018). Primarily geared towards the tokamak reactor type (figure 1), they include challenges across the board: plasma regimes of operation, the heat exhaust, neutron tolerance, tritium self-sufficiency, fusion safety, integrated design of a future power plant (DEMO) and economic viability. Finally, they also include research into the stellarator, a competitor to the tokamak, subject to many of the aforementioned challenges as well as some extra ones, but also with some marked advantages (Helander *et al*, 2012) Many of these challenges should be met by the International Thermonuclear Experimental Reactor (ITER) currently under construction in the South of France. A joint project by China, the European Union, India, Japan, Korea, Russia and the United States of America, its projected completion in 2025 should pave the way for answers to many—but not all—of the above while simultaneously breaking the $Q = 10$ barrier, the point at which the generated fusion power tenfold surpasses the input power.

The subject that underlies much of this thesis, divertor detachment, slots into that second mission: the heat exhaust. After the fuel has been 'burned' in the reactor vessel, after all, the power that was produced needs to be transported outside. The same goes for the 'exhaust fumes' that are produced along the way. The former is a necessary prerequisite for any power plant—that it produce power—and the latter, also not unintuitively, is necessary to prevent clogging of the reactor. The nature of a fusion reactor makes it such that the entire helium flux and a fixed fraction of the heat flux need to pass through a very narrow plasma zone called the scrape-off layer (SOL), which is consequently subject to an extremely high heat load. The SOL deposits all this power onto a likewise narrow region of the vessel wall, typically in the bottom of the reactor, called the divertor. The divertor cannot take the brunt of this heat flux—this is impossible from a materials perspective. An accepted value that appears throughout literature is a time-averaged limit of $10\ \mathrm{MW/m^2}$. This number appears throughout much of the literature related to divertors, although I could not trace it back to a specific publication. A recent paper by Pitts *et al* (2019) on the physics basis for the ITER tungsten divertor suggests that it might be as much as fifty percent higher. Be that as it may, there is still a huge gap between the maximum divertor heat flux and the $\mathrm{GW/m^2}$–order-of-magnitude heat flux that courses through the SOL. This heat flux needs to be reduced, and rather strongly at that.

By considering the units of heat flux, W over $\mathrm{m^2}$, two categories of mitigation strategies are possible: either one intervenes on the power side or one intervenes on the per-square-meter side. Successfully employed interventions exist on both sides. Impurity seeding, and consequent radiative cooling, has been used to radiate away a large part of the incoming power before it ever reaches the divertor target. Target tilting, meanwhile, has been used to act on the effective area that takes the remainder of the heat flux. By tilting the divertor target with respect to the incoming flux, i.e. positioning the surface normal of the divertor target at an angle to heat flux vector, the effective area that is exposed to this heat flux increases and thus the heat flux on the target drops.

Either technique has its limitations though. Unbridled impurity seeding has severe consequences for the performance of the reactor core, while the maximum tilt angle is restricted by the fact that the divertor

is not a continuous plane but rather consists of a discrete number of tiles of which the edges stand out ever so slightly, and which melt in the face of too strong target tilting (Pitts *et al*, 2017). Quantitatively, this means that the maximally realisable power drop due to impurity seeding is in the $80\%-90\%$ range. The minimum angle of incidence, meanwhile, is limited at $2.7°$ for ITER on the outboard side with shaping efforts to prevent leading-edge melting possibly increasing this value to over $4.2°$. These values are $3.2°$ and $4.7°$, respectively, on the inboard side (Pitts *et al*, 2019). The consequent reduction in heat flux, then, is somewhere in the $92\%-95\%$ range. Roughly speaking, either intervention independently yields a factor of ten reduction so combining the two yields a factor of a hundred reduction at best.

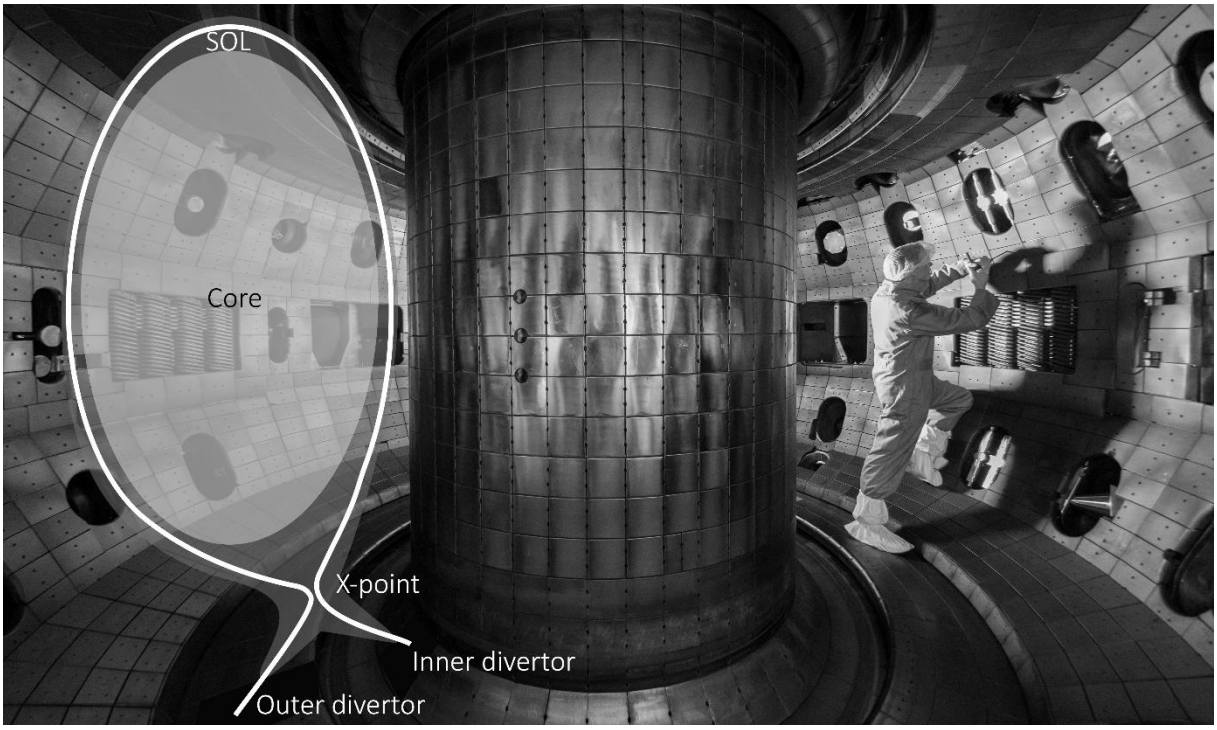

Figure 1.1 An inside view of the DIII-D tokamak in San Diego. Superimposed is a poloidal cross section of the plasma with the core (light grey), SOL (darker grey) and divertor areas (black) marked. Heat drifts out radially from the core until it hits the narrow SOL and courses along the SOL magnetic field (white line) to the inner and outer divertor targets. Adapted from: Rswilcox (2017).

For ITER—with a fusion power of $500\ \mathrm{MW}$ and a consequent SOL heat flux of about a $\mathrm{GW/m^2}$—these reductions combinedly bring the actual heat load within the operational window of the divertor, albeit only marginally so. For DEMO, a power-plant–scale device that is set to be built in ITER's wake, they will not suffice. To output electricity at power-plant level, Maisonnier *et al* (2005) estimate that the fusion power should be between $2.5\ \mathrm{GW}$ on the hopeful side, assuming advanced physics scenarios, and $5.0\ \mathrm{GW}$ on the conservative side, assuming limited extrapolation of existing physics and technology. Rising fusion power comes with an enhancement of the heat flux in the SOL— the footprint of the SOL does not scale at the same rate as the fusion power (Thornton, Kirk & MAST Team, 2014)—that cannot be countered via the measures proposed above, as those have already been employed to their fullest. It is here that detachment enters the equation.

Before moving into a discussion of detachment, though, this might be the moment to first develop in more detail some of the concepts above. Section 1.2 discusses some fusion fundamentals. Starting points are the fusion reaction that is at the heart of most fusion concepts nowadays, and the global power balance of a fusion reactor. Besides on the familiar plasma parameters of density and temperature, this balance turns out to be reliant on the energy confinement time. This leads up to a

discussion of magnetic confinement, tokamaks and stellarators. After covering the basics, a more detailed discussion of the core–SOL–divertor journey that heat and particles embark on is provided in section 1.3. Those familiar with nuclear fusion might instead decide to skip ahead directly to 1.4 and 1.5, which pick up the story of detachment and set out the research questions that this thesis should provide an answer to.

1.2 Basic principles of nuclear fusion

Nuclear fusion is the physical process in which two nuclei bind together to form one heavier nucleus in addition to possible by-products. Nuclear fusion reactions of which the product is not heavier than iron, which has the lowest binding energy per nuclear particle of all known elements, can in principle yield net energy but some of them are more effective than others. Mainstream fusion science is set on one particular reaction, that of deuterium (D) with tritium (T), which appears most favourable in the short run. This DT reaction is written as

$$\mathrm{D}^{+} + \mathrm{T}^{+} \rightarrow {}^{4}_{2}\mathrm{He}^{2+} + \mathrm{n} + 17.6\ \mathrm{MeV} \tag{1.1}$$

with the $17.6\ \mathrm{MeV}$ that is released during the fusion reaction divided unequally amongst both reaction products. This process has its optimal input power to product ratio at around $2 \times 10^{8}\ \mathrm{K}$ ($\sim 15\ \mathrm{keV}$). This temperature is so extremely high because the constituents of any fusion reaction have to approach each other to within an nuclear radius, which is where the strong nuclear force takes precedence over the coulombic repulsion of two positive charges. The threshold energy required to overcome the repulsive force up to that point is considerable; and would be higher still if not for the fact that ions have a finite chance of tunnelling through the coulombic potential barrier. Maintaining these temperatures without spending excessive power on heating equipment—which would defeat the purpose of energy generation—constitutes a major challenge to fusion. Crucially, it requires self-heating of the plasma (more on that later). A simplified way to capture the core dependencies of fusion is through the triple product, based on the Lawson (1957) criterion,

$$nT\tau_{\mathrm{E}} \geq \mathrm{constant}, \tag{1.2}$$

which gives a lower bound to the product of plasma density n, plasma temperature T and the energy confinement time $\tau_{\mathrm{E}} = W/P_{\mathrm{loss}}$, defined as the global plasma thermal energy W over the global power loss P_{loss}. In words, any fusion reactor must succeed in keeping enough particles at a high enough temperature, while insulating the system well enough. Note, in equation 1.2 it is assumed that the reaction rate of the fusion reaction scales with the temperature squared which limits its validity to the reactor-relevant range, somewhere between 100–200 million degrees kelvin (Lopes Cardozo, 2018).

Various solutions have been proposed to simultaneously solve these constraints over the years—some more viable than others—but the fusion community has since settled for *magnetic confinement,* because it is the concept that seems most promising in the short run. Put rather coarsely, magnetic confinement is 'just' an application of the Lorentz force. Although plasmas are neutral on large enough length scales, their individual constituents are charged and can thus be manipulated by the Lorentz force. This force causes any charged particle subject to a homogeneous magnetic field $\mathbf{B} = B_0\hat{\mathbf{z}}$ and with a velocity component perpendicular to that field into a circular orbit perpendicular to $\hat{\mathbf{z}}$, a process that is formally known as *gyration*, which cancels out any net movement perpendicular to the field.

Subjecting a plasma—which for the moment is assumed to be non-collisional—to a homogeneous magnetic field provides perfect particle isolation in any direction that sits perpendicularly to the magnetic axis $\hat{\mathbf{z}}$. It does nothing, however, to negate particle motion parallel to the field. And since no man-made field can stretch on *ad infinitum*, at some point the plasma will necessarily pour out of the ends.

The community thus found themselves at a junction. Either one tries to find a way to *cork* the ends of the field or one decides to do away with the ends entirely. A solution to the former was sought for decades, but try as they might, scientists never quite succeeded in corking their linear devices to the extent dictated by equation 1.2. A solution to the latter yielded better results. Mathematically, to *do away with the ends of the magnetic field* can be more precisely formulated as *finding a non-vanishing continuous vector field on an enclosed surface.* That it needs to be non-vanishing is crucial since a zero in the magnetic field would mean a hole in the insulation. That such a field should be continuous is likewise imperative, because any statement to the contrary would imply that $\nabla \cdot \mathbf{B} \neq 0$ somewhere. Finally, the statement that a magnetic field should sit on an enclosed surface means that, even should individual field lines never connect end-to-end, neither can they leave the enclosed surface, which in turn ensures that the plasma is locked onto such an enclosed surface. In the three-dimensional world that is our home, this rather limits options. For example, it rules out one go-to idealisation, the sphere, because it does not admit such a vector field on its surface. A relatively simple shape that does allow for such a vector field on its surface is the topological torus, three realisations of which are the frosted doughnut, the tokamak and the stellarator. Mapping the magnetic field onto such a torus, we can confine the plasma in all three dimensions.

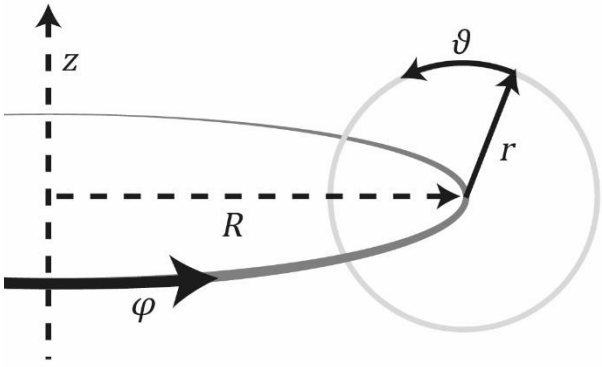

Figure 1.2. The coordinate directions that are used when discussing toroidally symmetric systems. They are the toroidal direction φ, the poloidal direction ϑ and the radial direction r. Together they form an orthogonal basis for a torus with a major radius R. The value for the major radius R is fixed while the value for r is bounded by the minor radius a for any given device. A third parameter that is sometimes given instead of the minor radius is the inverse aspect ratio $\varepsilon = a/R$ which is a scale-invariant combination of the minor and major radii. Fusion texts often refer to the *poloidal cross section* or *plane*, which refers to the plane spanned by ϑ and r. It is usually assumed that the magnetic field $\mathbf{B} = \mathbf{B}(r, \vartheta)$ only varies in this poloidal plane so that the plasma physics reduces to a two-dimensional problem, i.e. the toroidal direction φ is a symmetry direction.

Although this is certainly not the most intuitive way to build up the concept of toroidal confinement—it is often more simply described as a magnetic cylinder that, to suppress end losses, is connected end-to-end—it does result in plug-and-play closed flux surfaces: a fusion-technical term for *non-vanishing continuous vector fields on an enclosed surface*. A fusion torus is like a Matryoshka doll: to fill out the interior of a torus one needs a great many of these flux surfaces nested inside each other. Flux surfaces are critical to understanding transport, because heat and particles can (re)distribute freely and instantaneously on these surfaces, at least from the perspective of cross-field transport, and as such

plasma parameters like the density and the temperature are always constant on those surfaces. As a result a fusion torus becomes radially symmetric and core transport can be described purely by looking at its radial (r) component. Projected onto the poloidal plane (r, ϑ), these flux surfaces become flux tubes. The concept of a flux tube is highly relevant to this work. Such a flux tube needs not necessarily be closed, as is described in the next section.

Returning to the magnetic topology for one more moment, one last thing must be noted. Due to the slab-to-torus mapping a previously straight magnetic field now has curvature and gradients, which causes all manner of particle drift terms to emerge. One of them, the ∇B drift, was particularly problematic in the earliest days of fusion as it acts in opposite directions depending on the particle charge. The ensuing charge separation was detrimental to particle confinement. To solve this charge separation, a component of the magnetic field was introduced in the poloidal direction, ϑ in figure 1.2. This helical transform—reminiscent of the way a DNA string wraps around itself—connects the top and bottom of a fusion device along the magnetic field and thus short-circuits any build-up of charge.

The way in which this helical transform is introduced can be used to discriminate between today's two mainstream fusion concepts based on magnetic confinement: tokamaks and stellarators. Although both use coils external to the plasma for generating the toroidal (φ) magnetic field—these coils have been dubbed the toroidal field coils—they differ in how they impose the poloidal component. Tokamaks do so by using a central solenoid housed in the centre of the torus (the vertical dashed line in figure 1.2). By treating the central solenoid as the primary and the plasma as the secondary coil of a transformer and running a linearly ramped current through the former, an inductive plasma current can be driven in the toroidal direction. As prescribed by Ampère's circuital law, this plasma current will in turn induce a magnetic field in the poloidal plane thus achieving the helical transform. Alternatively, stellarators have this helicity built into their external coils *a priori*. The tokamak had more traction for a long time but nowadays it is realised that the stellarator has some marked advantages although it also has its share of additional issues (Helander *et al*, 2012). It is considered a serious contender to the tokamak in the long run.

1.3. Heat exhaust revisited: the core–SOL–divertor route

Before heat can be exhausted, it first needs to be generated and transported. Generation happens via the DT fusion reaction, equation 1.1, which primarily happens in the core of the plasma because it is hottest and densest there. After the reaction has taken place, there are two possible carriers of the energy generated: the alpha particle (synonymous to 'doubly charged helium ion') and the neutron. Marking the approximately fourfold weight of alpha particles compared to neutrons and considering conservation of momentum, it is readily found that 80% of the energy is imparted to the neutrons. From a plasma physicist's perspective, this heat instantly vanishes from the plasma. Or rather, the plasma never sees it in the first place. Neutrons do not interact with magnetic fields as they have no charge, and their cross sections are such that they do not stop for anything short of a 1-m solid wall (Freidberg, 2007). Granted, they do wreak havoc in that 1-m wall, but not because of the heat that they generate. A field that is likewise crucial to the future of fusion—it was enumerated as one of the key missions (EUROfusion, 2018)—neutron tolerance of materials is an area of ongoing research. However, it is not pertinent to this thesis.

The remaining 20% is for plasma physicists to contend with, though, as it is imparted to the alpha particles, which are influenced by the magnetic field. In an ITER-like scenario these alpha particles are

envisaged to be the primary heating mechanism. This marks a significant shift from current-day experiments, in which auxiliary heating mechanisms abound. In a quasi–steady-state scenario and assuming no auxiliary heating then, this 20% of the fusion power that goes towards heating has to be lost from the core at the same rate. Assuming that they do just that, once these alphas have thermalised with respect to the main plasma they are of no further use to the fusion effort. Or rather, they are very much unwanted as these 'helium ashes' clog the reactor core. It has been empirically determined that fusion reactors are subject to a limit in the plasma density, called the *Greenwald limit*, which caps the electron density (Greenwald, 2002). As helium ashes count towards this limit as well (and doubly as fast as hydrogen owing to their double charge), they reduce the maximum number of deuterium and tritium ions that are available for fusion. And since fusion is a two-particle process a loss of fuel density propagates quadratically into the generated power.

A pathway for the helium to leave the reactor has to be created in order to ameliorate these issues. In order to do so, the flow of the plasma needs to be diverted (hence the name 'divertor'). Diverting the flow of the plasma in turn requires diverting the direction of the magnetic field. To this end a *divertor coil* is included in the bottom of the device, which generates an additional poloidal field such that a magnetic null is created in the poloidal plane between the magnetic axis of the plasma core (at zero minor radius, $r = 0$) and the magnetic axis of the divertor coil. This position is called the X-point. By inspecting figure 1.3 one easily sees why. The flux surface that sits directly above the X-point towards the reactor core is called the last closed flux surface (LCFS) and any flux surface that sits inside that is part of the core plasma, while any flux surface that sits below the X-point is part of the private flux area of the divertor. Moving to the left or right of this X-point, flux surfaces are formed that must now round both the core and the private flux area, because they cannot cross the magnetic null nor intersect the flux surfaces of the private flux area. However, rather than allowing the plasma to complete its new, dumbbell-shaped trajectory a pair of divertor plates are added on which these trajectories terminate. This configuration is termed the single-null configuration, referring to the formation of one X-point. Alternatives such as the double-null configuration also exist, which have a similar divertor construction at the top resulting in two positions of zero magnetic field, c.q. two X-points.

The region of the plasma that sits outside of the LCFS is called the scrape-off layer. However, even in speaking, it is usually referred to as the SOL. The SOL forms the outermost layer of the plasma and collects the heat and particles that come out of the hot plasma core. It is exceedingly thin in the radial direction, $d_{\mathrm{SOL}} = 3.4 \pm 0.5\ \mathrm{mm}$, in ITER when averaged over many SOLPS-ITER runs. However, other techniques yield results ranging from $1\ \mathrm{mm}$ to $6\ \mathrm{mm}$ (Pitts *et al*, 2019). Taking d_{SOL} as that value, since it sits somewhat in the middle, one finds the toroidal footprint (the SOL thickness times the major radius circumference) to be very small, $A_{\mathrm{SOL}} \approx 2\pi R_{\mathrm{ITER}} d_{\mathrm{SOL}} = 0.13\ \mathrm{m}^2$. That area is equal to this page and the next printed in A4.

All the alpha particle heat that comes out of the fusion core does so radially (cross-field) until it passes the LCFS; see figure 1.4. When it reaches the SOL it is instead transported predominantly poloidally, because parallel transport along the magnetic field is many orders of magnitude faster than cross-field transport. The SOL collects all this heat, and transports it through its very thin layer. This means that the resulting parallel heat flux in the SOL for an ITER-like scenario with $100\ \mathrm{MW}$ of alpha heating is on the order of $1\ \mathrm{GW/m^2}$. Complicating matters is the fact that d_{SOL} appears mostly unaffected by machine size, which means that a tenfold increase in power production moving from ITER to DEMO does not go hand in hand with a tenfold increase in the SOL footprint. For DEMO, this means that SOL loading will significantly exceed $1\ \mathrm{GW/m^2}$. Target tilting and impurity seeding, mitigation techniques that stretched to their limits can just about keep divertor power loads in ITER in check, are thus not sufficient for DEMO.

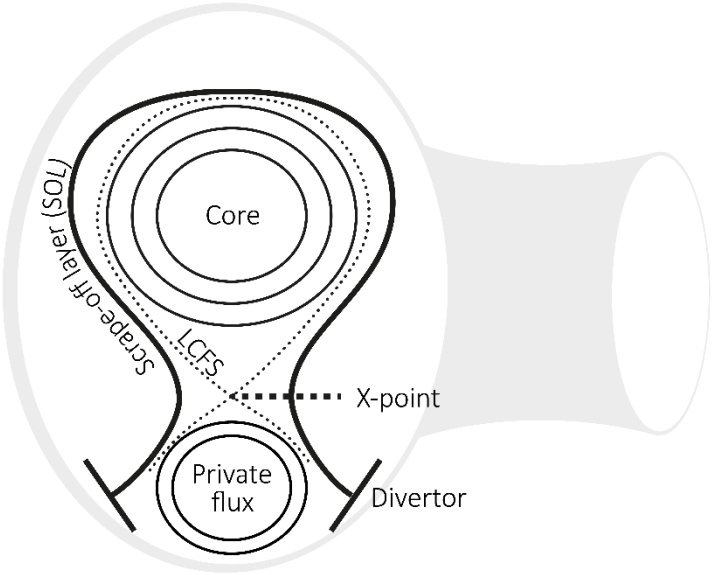

Figure 1.3. Poloidal cross section of a torus with a single-null configuration. Using a divertor coil an additional poloidal magnetic field can be constructed in the same direction as that of the main plasma, such that they cancel somewhere in between. This position of zero magnetic field is called the X-point. The flux surface directly above this is called the last closed flux surface (LCFS), with any flux surface within the LCFS part of the core plasma, while any flux surface that sits under it is part of the private flux area. Any flux surface that sits outside the LCFS must round both the core plasma and the private flux region. However, two divertor plates are in the way. The area of the plasma outside of the LCFS, in which field lines terminate on the divertor plates, is called the scrape-off layer (SOL).

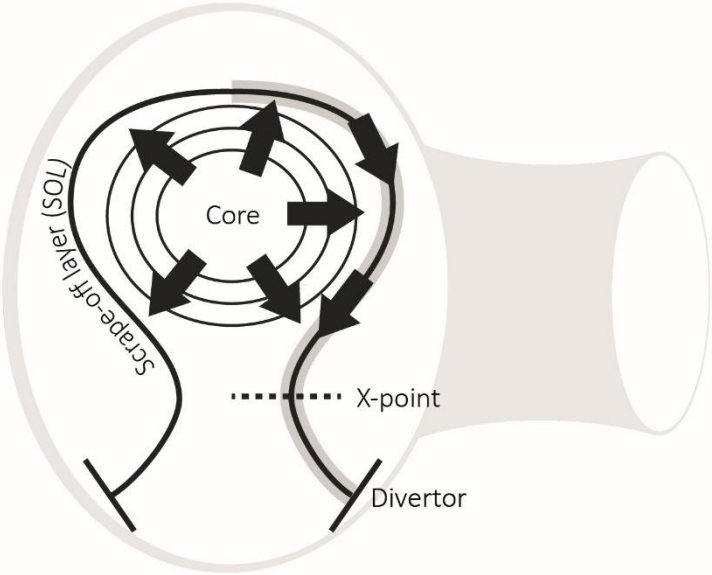

Figure 1.4. This poloidal cross section show the journey of heat and particles through a fusion device. DT reactions take place in the hot plasma core. After their conception, both the generated power and the helium 'exhaust fumes' start moving radially outward. This happens in a radially symmetric fashion because both pressures and temperature are in equilibrium on the flux surfaces (the thin-lined circles). At a certain position from the centre, heat and particles pass the last closed flux surface (LCFS, not drawn). After this LCFS, heat and particles enter the scrape-off layer region of the plasma. What sets this region apart from the core region is that field lines, rather than staying on enclosed flux surfaces, now 'terminate' in the bottom of the device. Since transport along the magnetic field is so much faster than transport across it, the scrape-off layer is depleted extremely fast and thus very thin. Whosoever quirked an eyebrow at the statement that field lines terminate, is fully right. They cannot truly terminate as this would imply $\nabla \cdot \mathbf{B} \neq 0$; rather, they continue along a dumbbell-shaped trajectory of which the point of closest contact is the X-point. However, heat and particles cannot follow there as there is a solid (or possibly liquid) wall in the way: the divertor.

1.4. Divertor detachment and the motivations for this thesis

To keep the divertor within operational limits requires something extra on top of the previously discussed impurity seeding and target tilting mechanisms: divertor detachment. At a SOL level, the detached state is one three possible states of the plasma. The others are the sheath-limited and conduction-limited (high-recycling) regimes. Stangeby (2000) calls these the simple and the complex SOLs respectively. The former assumes a constant temperature across the SOL length and no sources or sinks of particles or energy except on the target, while the latter allows for parallel gradients as well as volumetric radiation losses, momentum losses and particle losses between the upstream and target sides. Operating the SOL in the sheath-limited regime is ill-advised as it means that the divertor temperature, which is low, would match the upstream temperature, and that should ideally be much higher. The conduction-limited region does accommodate such a temperature drop across the SOL. The detached regime, finally, can be transitioned into from the conduction-limited regime. There is, however, no universally accepted definition of detachment so the point at which the SOL ceases to be conduction-limited and becomes detached instead is somewhat arbitrary. Stangeby (2018) defines detachment as the point at which the temperature at the divertor drops below $10\ \mathrm{eV}$ ($\sim 1.2 \times 10^5\ \mathrm{K}$). Note that this definition holds on a per-flux-tube basis, so any given flux tube in the SOL can be in the detached regime while others are not. ITER is envisaged to operate in a mode of partial detachment in which the strike point, i.e. the part of the divertor that takes the highest heat load, is detached but the rest of the divertor is not (ITER Physics Expert Group on Divertor *et al*, 1999).

1D modelling efforts (Dudson *et al*, 2019) have shown that two minimal requirements for detachment to occur are (1) a loss of power and (2) a loss of momentum along the SOL. Both of these factors, it should be noted, contribute to why detachment is so helpful to the nuclear fusion effort. As discussed previously, a prime motivator for moving into the detached regime is the necessity to shed sufficient power before the divertor target. Meanwhile, the level of erosion of the divertor target is linearly proportional with the particle flux onto the target, which is influenced in turn by the level of momentum losses. There is also a large power deposition associated with this, as all the energy due to recombination is absorbed by the divertor surface (Stangeby, 2000). If either of the two is missing the other requirement cannot cause the SOL to enter the detached regime. The onset of detachment thus cannot be attributed to any single physical process. Rather, there are a whole bunch processes that factor in ranging from ionisation and excitation to impurity radiation and recombination. Understanding what processes affect detachment to what extent is crucial to tweaking the detached regime.

Experiments in e.g. JET (Field *et al*, 2017) and TCV (Harrison *et al*, 2017) indicate that the detached regime is also subject to temporal evolution. Today's workhorse scrape-off layer modelling suite, SOLPS-ITER, is a dynamics solver that solves the full transport equations for all present species in a poloidal geometry (Schneider *et al*, 2006) in principle. However, the complexity and computational resources required for time-dependent modelling in SOLPS-ITER preclude its usage for extensive dynamics simulations (Bonnin *et al*, 2017). There are opportunities for simpler models to describe the dynamics of the scrape-off layer based on a strongly simplified physics basis. One such model, DIV1D, is an in-house Fortran code developed by Westerhof (2019) at the Dutch Institute For Fundamental Energy Research (DIFFER). It closely resembles SD1D (Dudson, 2016; Dudson *et al*, 2019), a C++ variant based on much the same physics and numerics.

This DIV1D code was considered a useful tool to study the dynamics of divertor detachment. The whole of this dynamics is still a rather broad subject, so in particular this project was to focus on the way that radiation fronts (positions of peak intensity for some of the key physical processes related to

detachment as described above) behave in the face of two transient phenomena: gas puffs and edge-localised modes (ELMs). The former is an externally controlled effect—gas puff is a fuelling mechanism for supplying the fuel that a fusion reactor needs—while the latter is contingent on the plasma regime and whether or not the instabilities that are the root cause of ELM formation exist or not (Hill, 1997). Gas puff acts on the particle balance of neutrals (and thus via ionisation on that of the plasma) while ELMs acts on the heat flux and the particle balance.

1.5. Research questions

Initially, the plan for this project was to focus exclusively on the effects of those two effects—gas puff and ELMs—on the dynamical behaviour of the detachment front. However, over the course of several months these research goals evolved. There were two main stages to this evolution. The first of them took root over the course of several discussions with T. Ravensbergen and J. Koenders at DIFFER, who were at the time working on the two-point model for control purposes. To this end, they drew on SOLPS-ITER simulations for the Tokamak à Configuration Variable (TCV), which were devised and run by M. Wensing (2020) at EPFL. We decided to try if we could not reproduce those results from SOLPS-ITER in DIV1D, which resulted in to a rough preliminary research question: Are the results from DIV1D in accordance with those from SOLPS-ITER?

The answer was no. Within the available parameter space of DIV1D no results could be obtained that were as much as qualitatively similar. We will not dive into the peculiars, here, because that would require taking an advance on the subject matter of chapters 2 and 4. It brooked no argument, however, that this discrepancy between SOLPS-ITER and DIV1D was one that needed to be understood in more detail. After all, there are opportunities for simpler 1D models, but that only holds insofar as those simpler models are still able to accurately capture the physical reality. One might interject at this point that SOLPS-ITER is still 'just' a modelling suite, and not the physical reality. One would be correct in stressing that. However, SOLPS-ITER is also a design tool used for estimating the operational parameters of future devices and experiments cannot probe what has yet to be built. So at the very least, it is one of the better estimates at that future physics. Going back to the matter at hand, we were faced with a code that should but would not be corroborated by SOLPS-ITER outputs while we had hoped to obtain at least qualitatively similar results. This motivated the eventual formulation of a more refined research question:

> What physics mechanisms drive the differences between DIV1D and SOLPS-ITER, and can they be accounted for in 1D?

To answer this research question required understanding several aspects. These aspects occurred sequentially rather than in parallel, and by discussing them now we run the risk of giving away the gist of the story. As this is a thesis and not a novel, however, we will have to accept that. First and foremost, it required a good understanding of the plasma physics incorporated in the DIV1D code. Any hopes of achieving a similar understanding for SOLPS-ITER proved to be in vain. It quickly turned out that a ground-up approach—while perfectly feasible for DIV1D—was ill-advised for a code whose documentation numbered in the hundreds of pages with some formulae requiring the better part of a page. Initial attempts to neatly reconstruct the power balance of SOLPS-ITER did not work out as planned either, and were eventually discarded. What this early phase did suggest, though, was that power losses perpendicular to the flux tube might be a significant factor: the radial losses. These are present in SOLPS-ITER but not in DIV1D. Taking a more phenomenological tack, then, we simply

implemented several types of perpendicular losses in DIV1D to see what their effects on the key plasma parameters[1] would be.

A small note on the terminology with respect to the terms 'radial' and 'perpendicular' is in order here. There are two directions perpendicular to the flux tube. One exists in the plane of the flux surface and has components in ϑ and φ—let us call it the in-plane perpendicular direction—while the other is perpendicular to both the flux tube and the flux surface and has a component solely in r—let us call it the radial direction—with ϑ, φ and r as in figure 1.2. Since the in-plane perpendicular direction is a symmetry direction there is no loss channel in that direction. The only perpendicular losses that could possibly be incurred are perpendicular losses in the radial direction. Because of this the terms perpendicular and radial are used interchangeably in this work with the understanding that they always refer to the direction perpendicular to the flux surface.

Moving to the other research goal of this graduation project, we also ended up doing things a bit differently than we envisaged at the onset. While the idea was to investigate both the effects of gas puffs and ELMs, the current work is limited to only a discussion of the effects of ELMs. While a simple gas puff model was implemented in the DIV1D code over the course of this project by Egbert Westerhof, preliminary simulations in which that simple gas puff was perturbed by a sinusoidal fluctuation showed an instantaneous and rather trivial response of the plasma. Additionally, experimental data analyses by Koenders (2020) would later show that the gas-puff–dependent dynamics of the plasma were predominantly determined by the dynamics of the gas valve alongside a diffusion equation, both effects with no straightforward link to the current state of the DIV1D code. This left us with very few leads. Focussing, rather, on the dynamics of the plasma due to ELM perturbation, we formulated the following research question:

> What is the effect of ELMs on the dynamical behaviour of radiation fronts in the scrape-off layer?

To begin this begged the question of how to implement ELMs in DIV1D. That started with identifying the key features of an ELM. We had to determine what key features a simple ELM model needed to incorporate, which parameters that involved and what the typical sizes of those parameters were. These parameters in hand, we could then attempt to program different types of ELMs into DIV1D to determine which yielded the best results. This feature implemented, we could embark on a characterisation of the dynamics. This involved tracking the detachment front, and determining what happened to it—not only as a function of time but also as a function of for example the upstream density. Closely linked to this, it allowed us to determine whether an ELM caused reattachment or not. Also of interest was the added integrated target heat flux (the fluence) due to an ELM, and how it was influenced by ELM amplitude and frequency.

1.6. Reading guide

The body of this thesis is split up in four chapters. In chapter 2, titled *Modelling the scrape-off layer*, we start out from the general plasma fluid equations in section 2.1 and peel off dimensions one by one in ensuing sections. Section 2.2 briefly sketches the two-dimensional representation of the SOL used in SOLPS-ITER, lingering at the role of perpendicular losses due to anomalous transport in particular.

[1] A discussion of what these key plasma parameters are is left for the next chapter.

Section 2.3 extensively discusses one-dimensional DIV1D. The 1D fluid equations that govern it are derived from the general plasma fluid equations, the transport equation for the neutrals is presented and the different sources and sinks are elaborated on. Section 2.4, finally, discusses an analytical solution to the 1D SOL given several simplifying assumptions: the two-point model. Besides being an excellent way to discuss the key SOL plasma parameters and their relations to one another, it segues nicely into a very basic discussion of how detachment enters into the equation.

Chapter 3, titled *Code verification*, is dedicated to determining whether the code works as it should. Section 3.1 discusses the results of a grid convergence scan. This is primarily to show that the spatial resolutions at which this thesis' simulations were run are sufficient for present purposes, but also to get a sense of the size of any residual uncertainty. Section 3.2 discusses the results of scanning the artificial viscosity. Although instrumental to make the code 'act nice' it—as its name suggests—is an unphysical artifice, which makes it important to determine to what extent it influences results.

Chapter 4 is titled *Radial power losses.* It is built up in a chronological fashion, starting with a description of a specific SOLPS-ITER shot (Wensing, 2020) that would not let itself be captured by DIV1D physics alone in section 4.1. It also delivers a set of DIV1D-compatible input parameters that serve as the starting point to the simulations in section 4.2. By tweaking DIV1D's parameters one by one, that section shows what can and cannot be influenced by these parameters. By the end of that chapter the reader will understand that the physics features of DIV1D that existed at the time did not suffice to describe the results found in SOLPS-ITER. Section 4.3 builds on this understanding and discusses the attempt to amend this by—as the chapter title already suggested—introducing radial losses to the power balance.

Chapter 5 is titled *ELM-induced dynamics.* Section 5.1 covers some basic aspects of ELMs and discusses how they can be used to develop simple ELM models for use in DIV1D. One model is discussed from literature, and a self-defined alternative to that model is proposed. Having covered the implementation of ELMs in DIV1D, we put them to the test in section 5.2 by making a case study of ELMs in a MAST like divertor leg. This section discusses back-of-the-envelope calculations to get a sense of the relevant parameter space, and discusses several parameter scans that cover the main aspects of the ELM models and the influence that they have on the divertor plasma.

Chapter 6, titled *Summary and discussion*, recapitulates the main findings of chapters 3 through 5 and, insofar as they have not been discussed sufficiently in the chapters themselves, adds to their discussion. Finally, chapter 7, titled *Conclusions and outlook*, gives concise answers to the two research questions of section 1.5 and closes off with some potentially interesting tangents for further research.

2. Modelling the scrape-off layer

Modelling the plasma fluid behaviour in the scrape-off layer is complex work. Strategies to cope with such complexity abound. They take on various forms for myriad reasons. Such things as the plasma parameters of interest, the timescales and length scales that are involved, a wish for statics or dynamics calculations and more mundane factors like the available amount of computational power, inform decisions about what constitutes *the best model for the job*. It is the main purpose of this chapter to show that DIV1D—a one-dimensional scrape-off layer code that simulates the part of a fusion plasma that lives between the X-point and the divertor target—provides an interesting framework to bring to light some, if not all, crucial aspects of divertor detachment.

This chapter reads not unlike the way that one peels layers off an union. As a rather standard but nonetheless crucially important starting point, the general plasma fluid equations will be treated. The unpeeled union, so to speak, not yet restricted by any geometry. In the bulk of this chapter progressively fewer-dimensional geometries are imposed. They are exemplified by three models: SOLPS-ITER, DIV1D and the two-point model (2PM). Forgoing a three-dimensional treatment because it is of no direct consequence to this thesis (which does not mean that it is of no consequence), these are the two- and one-dimensional models. Finally, an analytical solution of a strongly simplified one-dimensional model is presented. Despite being wedged in the middle, a discussion of DIV1D is the centrepiece of this chapter. It includes derivations of the DIV1D fluid equations from the general plasma fluid equations, a summary of the atomic (collisional) physics taken into account and the treatment of the neutral gas.

SOLPS-ITER and the 2PM are important ingredients along the way. SOLPS-ITER will serve as the basis for comparison when dealing with the question of DIV1D's statics. It is informative to have a bird's-eye view of what that code does. What happens under the hood of SOLPS-ITER, however, will for the bigger part stay there as understanding the results of this thesis requires little understanding of how SOLPS-ITER actually work. The 2PM, on the other hand, will be zoomed into in more detail. It is systematically stripped of almost all physics and thus well-suited to introducing the key plasma parameters that are involved in modelling the scrape-off layer. Moreover, using a few straightforward corrections it can be adapted to allow for divertor detachment as well. This makes it a perfect candidate for honing one's physical intuition regarding the scrape-off layer and detachment physics.

2.1. General plasma fluid equations

This section closely follows section 2.3 of Lieberman & Lichtenberg (2005). Within a six-dimensional phase space $(\mathbf{r}, \mathbf{v})$, a distribution function $f(\mathbf{r}, \mathbf{v}, t)$ can be defined on a per-species basis[2] such that $f(\mathbf{r}, \mathbf{v}, t)\mathrm{d}^3r\mathrm{d}^3v$ contains all particles of that species in a volume $\mathrm{d}^3r\mathrm{d}^3v$ centred at $(\mathbf{r}, \mathbf{v})$ at time t. A continuity equation can then be derived for $f(\mathbf{r}, \mathbf{v}, t)$, which takes the shape

$$\frac{\partial f}{\partial t} + \mathbf{v} \cdot \boldsymbol{\nabla}_{\mathbf{r}} f + \mathbf{a} \cdot \boldsymbol{\nabla}_{\mathbf{v}} f = 0 \tag{2.1}$$

[2] The distribution function as well as derived quantities, like a species' density n, should really be labelled f_s and n_s since these variables exist for each species s. However, as the present derivation is independent of s anyway, those subscripts have been dropped in favour of a simpler notation.

where the gradient operators read $\boldsymbol{\nabla}_{\mathbf{r}} = \hat{\mathbf{x}}\partial/\partial x + \hat{\mathbf{y}}\partial/\partial y + \hat{\mathbf{z}}\partial/\partial z$ and $\boldsymbol{\nabla}_{\mathbf{v}} = \hat{\mathbf{x}}\partial/\partial v_x + \hat{\mathbf{y}}\partial/\partial v_y + \hat{\mathbf{z}}\partial/\partial v_z$. Equation 2.1 is called the *collisionless Boltzmann equation* or *Vlasov equation*. The modifier *collisionless* stems from the fact that, in deriving equation (2.1), the only changes to $f(\mathbf{r}, \mathbf{v}, t)$ are effected by particles entering via the faces of $\mathrm{d}^3 r \mathrm{d}^3 v$. Nothing happens within the volume itself. In plasma reality, though, all manner of interparticle collisions may exist: interactions that create additional particles of one species while removing those of another.

To account for these effects, it is necessary to expand on equation 2.1 by adding a collision term. These interparticle collisions are assumed to occur on a much shorter timescale than that of the evolution of $f(\mathbf{r}, \mathbf{v}, t)$ itself, and can thus be considered to effect instantaneous changes to the velocity of a particle. By adding the collision term, which accounts both for addition (sources) and removal (sinks), on the right-hand-side of 2.1, one obtains the *Boltzmann equation*

$$\frac{\partial f}{\partial t} + \mathbf{v} \cdot \boldsymbol{\nabla}_{\mathbf{r}} f + \mathbf{a} \cdot \boldsymbol{\nabla}_{\mathbf{v}} f = \left.\frac{\partial f}{\partial t}\right|_{\mathrm{c}}. \tag{2.2}$$

The "usual" plasma fluid equations follow from these equations by calculating the first three moments of $f(\mathbf{r}, \mathbf{v}, t)$ with respect to $\mathbf{v}$, which yields the density $n(\mathbf{r}, t)$, the particle flux $\Gamma(\mathbf{r}, \mathrm{t})$ and the kinetic energy per unit volume $w(\mathbf{r}, t)$ as

$$n(\mathbf{r}, t) = \int f \mathrm{d}^3 v\,; \tag{2.3}$$

$$\Gamma(\mathbf{r}, t) = n\mathbf{u} = \int \mathbf{v} f \mathrm{d}^3 v\,; \tag{2.4}$$

$$w(\mathbf{r}, t) = \frac{3}{2} p + \frac{1}{2} m n u^2 = \frac{1}{2} m \int v^2 f \mathrm{d}^3 v. \tag{2.5}$$

Here m denotes the species' mass, and $\mathbf{u}(\mathbf{r}, t)$ the location- and time-dependent average velocity, while $(3/2)p$ and $mnu^2/2$ should be read as the static and dynamic components of the kinetic energy per unit volume. Conservation laws can be found for all three parameters by appropriate integration of equation 2.2. This yields the familiar plasma fluid equations.

Firstly, the particle balance is given by

$$\frac{\partial n}{\partial t} + \boldsymbol{\nabla} \cdot (n\mathbf{u}) = S_{\mathrm{density}}\,, \tag{2.6}$$

where S_{density} is the sum of all particle sources and sinks that results from integrating the collision term on the right-hand side of equation 2.2.

Secondly, the momentum balance is given by

$$mn\left[\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \boldsymbol{\nabla})\mathbf{u}\right] = qn(\mathbf{E} + \mathbf{u} \times \mathbf{B}) - \boldsymbol{\nabla} p + S_{\mathrm{momentum}}, \tag{2.7}$$

which is valid when assuming an isotropic pressure. More generally, $\boldsymbol{\nabla} p = \boldsymbol{\nabla} \cdot \boldsymbol{\Pi}$, where $\boldsymbol{\Pi}$ is the pressure tensor. From left to right, these terms are due to a plasma acceleration, an inertial effect, an

electric field (with q the species' charge), a magnetic field, a pressure gradient, and some combination of collision terms that are lumped in S_{momentum}. In many cold plasma scenarios it is common to drop the magnetic term immediately—as is done in Lieberman & Lichtenberg—because low temperatures combined with high densities effect a highly collisional plasma in which mean free paths are so short that electrons and ions cannot follow along a magnetic field line for any appreciable amount of time. In the context of fusion plasmas, however, this term is decidedly non-negligible and it is thus retained.

Thirdly, the energy balance is given by

$$\frac{\partial}{\partial t}\left(\frac{3}{2}p\right) + \boldsymbol{\nabla}\cdot\frac{3}{2}(p\mathbf{u}) + p\boldsymbol{\nabla}\cdot\mathbf{u} + \boldsymbol{\nabla}\cdot\mathbf{q} = Q. \tag{2.8}$$

From left to right, these terms represent the evolution of the thermal energy, the convective heat flux, heating or cooling due to compression or expansion, the conductive heat flux and some combination of collisional sources and sinks that are lumped in Q.

Now, assuming $\mathbf{E}$ and $\mathbf{B}$ to be externally controlled parameter while simultaneously assuming that the source terms S_{density}, S_{momentum} and Q can be described in terms of one or more parameters that exist within the equations, equations 2.6–2.8 still do not constitute a closed set. Two further expressions are required, one for the (partial) pressure p and one for the conductive heat flux $\mathbf{q}$. For the pressure, one possibility is to use the following equation of state

$$p = nk_{\mathrm{B}}T, \tag{2.9}$$

where k_{B} is the Boltzmann constant and T is the temperature in K. For the conductive heat flux, one possibility is Fourier's law (Fourier, 1822), which defines $\mathbf{q}$ in terms of the thermal conductivity κ and a gradient of the temperature following

$$\mathbf{q} = -\kappa\nabla T. \tag{2.10}$$

2.2. Two dimensions: SOLPS-ITER

2.2.1. B2.5 and Eirene

SOLPS-ITER—with SOLPS (read SOL–P–S) an acronym for Scrape-Off Layer Plasma Simulator—is a code package that simulates the SOL on the poloidal plane. It combines B2.5, a fluid code that is used to simulate the charged species subject to the magnetic field, with Eirene, a Monte-Carlo–based kinetic code used to simulate the neutral species. SOLPS has existed for quite a few years and its underlying coupling between B2 (a predecessor to B2.5) and Eirene for even more, since the early 1990s. The suffix -ITER reflects that the ITER Organization has since realised the importance of SOLPS-ITER to its modelling efforts and has taken an active part in updating SOLPS to the state-of-the-art, culminating in the release and distribution of SOLPS-ITER in 2015 (Bonnin & Pitts, 2015).

B2.5 solves the Braginskii equations (Braginskii, 1965; Schneider *et al*, 2006). Like the treatment above, the Braginskii equations are derived from the Boltzmann equation. As a result, the set of equations that results is in principle not unlike equations 2.6–2.8, which are fully general except for the omission the

pressure tensor. However, the righthand sides—atomic sources and sinks, interaction terms between different species and effects such as viscosity—are all lumped together in the S's. The Braginskii equations, although physics wise similar, have these terms explicitly developed in much greater detail. These equations are solved on a field-aligned numerical grid. The field-aligned direction (or to be precise, the projection of the field-aligned direction onto the poloidal plane) is given by x while the perpendicular component is given by y. Because cell sizes differ from position to position and because the pitch angle ($\sin(\vartheta) = B_\vartheta/|\mathbf{B}|$) varies over the poloidal plane, the B2.5 equations include a host of geometric correction factors that make them look rather unintuitive to a novice.

Adding further complexity to the code is that these equations need to be solved for not just one or two but a host of species. Even in the relatively straightforward TCV scenario taking into account a deuterium plasma contaminated with just carbon impurities, as is the case in the SOLPS-ITER simulation used in this thesis, this requires tracking one ionised hydrogen species and six ionised carbon species (one for each charge state), all of which couple with each other and with the neutral hydrogen and carbon atoms. The behaviour of these neutrals cannot generally be described by B2.5, although it can in certain situations be used to. In such a case the neutrals are treated as any other plasma species, but omitting those terms that find their origin in the charged nature of species. However, because B2.5 is a fluids code while the behaviour of neutrals at higher temperatures, $T > 10^5\ \mathrm{K}$ (or $10\ \mathrm{eV}$), is kinetic in nature (Schneider *et al*, 2006), this is not generally an option. In SOLPS-ITER B2.5 is consequently coupled to Eirene, which is a Monte-Carlo–based kinetic solver. Using Eirene incurs significant run-time penalties, however, as it produces statistical results that require extra care to attain convergence.

A detailed discussion of all the possible physical effects that exist in SOLPS-ITER would do little in terms of clarification. Taking a different tack entirely, the physics that is relevant for the SOLPS-ITER–DIV1D comparison will be introduced on the fly in chapter 3. For a complete reference to the relevant physics, the interested reader is directed instead to Schneider *et al* (2006).

2.2.2. Anomalous transport

One issue that does require discussion is that of anomalous transport as it strongly factors into the discussion of radial losses in SOLPS-ITER and DIV1D later. Under the assumptions of classical transport and of axisymmetry, i.e. symmetry with respect to the toroidal angle φ, a torus becomes an intrinsically two-dimensional system. In such a case a poloidal projection like in SOLPS-ITER is without consequence. Tokamaks depart from axisymmetry only slightly—on account of a discrete number of poloidal fields coils slight variations in the toroidal magnetic field exist—while stellarators do so strongly (Boozer, 2018). The first assumption, that of classical transport, is quite a bit more restrictive. Early prognoses for transport levels in tokamaks were always wide of the mark, underestimating transport coefficients by several orders of magnitude. As a result, a new field emerged: that of neoclassical transport. Neoclassical transport ensues when the toroidal direction, previously neglected, is taken into consideration. The resulting transport can eclipse that of a classical treatment by two orders of magnitude (Freidberg, 2007). Yet the story does not end there. In spite of the inclusion of neoclassical effects, discrepancies persisted which led people to conclude the existence of some form of 'anomalous transport', or more colloquially, turbulence. As Schneider *et al* (2006) point out:

> "Despite the fact that SOL modelling is very successfully able to identify basic physics mechanisms and represent even quantitatively the experimental results, the largest uncertainty in the SOL description is the question of anomalous transport."

Turbulence is not covered in SOLPS-ITER physics basis, though, and is thus a major driver for uncertainty. To have the code better conform to experimental reality a set of anomalous transport coefficients is introduced for the radial particle diffusion, and the radial ion and electron heat diffusion coefficients. These are set by the modeller him or herself and strongly influence results. An effort to couple SOLPS-ITER to BOUT++ is currently underway (Zhang *et al*, 2019; Zhang *et al*, 2020) to model exactly these parameters in a self-consistent fashion, which further underpins their importance. This is not the first such effort. BOUT++ is a three-dimensional turbulence code whose predecessor BOUT was previously coupled to UEDGE, also a scrape-off layer code, with a self-consistent model in mind (Rognlien *et al*, 2004). The existence of these coefficients might be a key factor in why the one-dimensional code that is at the heart of this thesis so poorly reproduced SOLPS-ITER results at the onset of the project. The results and discussion related to this are the subject matter of chapter 3.

2.3. One dimension: DIV1D

In the previous section SOLPS-ITER was introduced as a way of modelling the scrape-off layer in two dimensions. Attempts to model the scrape-off layer in this fashion are met with one considerable disadvantage. They require hefty amounts of computational resources to compute steady-state solutions. Should one wish to do dynamics simulations with SOLPS-ITER, one faces even larger computational loads. In the worst case, the computational load scales with the number of points in time. Even if measures were introduced to make this scaling more computationally favourable, one is still left with a load that starts out heavy and keeps becoming heavier with increasing temporal resolution. There are opportunities, then, for faster alternatives—at least if those alternatives perform well enough. This section introduces one such alternative, DIV1D. Figure 2.2 on page 27 schematically shows what it means for the model to become 1D.

2.3.1. Derivation of the 1D plasma fluid equations

In this section, starting out from the equations 2.6–2.10, the DIV1D equations are derived, which are simply the 1D plasma fluid equations under the assumptions of equal electron and ion densities, and thermal equilibrium between both species. The domain of DIV1D, like is the case in SOLPS-ITER for the x component, is defined parallel to the magnetic field.

The DIV1D continuity equation is identical to equation 2.6, albeit in one dimension. Because it is assumed *a priori* that the electron and ion densities are equal just one continuity equation is required that solves for both the electrons and ions, which is given by

$$\frac{\partial n}{\partial t} + \frac{\partial \Gamma_n}{\partial x} = S_{\text{density}}. \tag{2.11}$$

Here $\Gamma_n = nv_{\parallel}$ is the particle flux with $v_{\parallel}$ the velocity along the magnetic field (which is parallel to the domain of DIV1D).

Box 2.1. Simplifying assumptions to the DIV1D model

Quasi-neutrality dictates that the electron density n_e and the ion density n_i should be equal in case of hydrogen isotopes. It is assumed that electrons and ions are in thermal equilibrium, which means that their respective temperatures T_e and T_i are equal. Furthermore, in 1D the gradient operator, the plasma flow velocity $\mathbf{u}$ and the magnetic and electric field vectors, $\mathbf{B}$ and $\mathbf{E}$, simplify to only the component parallel to the magnetic field, denoted by subscript $\|$. Because charges can move freely along a field line, charge separation cannot occur and hence the parallel electric field is always zero.

Summarised formulaically, the simplifications employed by DIV1D with respect to the general plasma fluid equations 2.6–2.10 are

(I) $n_\mathrm{e} = n_\mathrm{i} = n;$ (II) $T_\mathrm{e} = T_\mathrm{i} = T;$ (III) $\nabla = \hat{\mathbf{x}}\partial/\partial x;$

(IV) $\mathbf{u} = v_\| \hat{\mathbf{x}};$ (V) $\mathbf{B} = B_\| \hat{\mathbf{x}};$ (VI) $\mathbf{E} = E_\| \hat{\mathbf{x}} = \mathbf{0}.$

The DIV1D momentum balance simplifies considerably compared to equation 2.7. Due to the field alignment of DIV1D, the effect of any electric field would be immediately cancelled since particles can move freely along the field, so $E_\| = 0$. Since the flow velocity is along the direction of the magnetic field, the magnetic term likewise evaluates to zero. After some additional manipulation (involving the chain rule, then inserting equation 2.11 and finally absorbing the particle source term into the lumped momentum source term S_momentum) it is found that

$$\frac{\partial (m n v_\|)}{\partial t} = -\frac{\partial}{\partial x}\left(m n v_\|^2 + p\right) + S_\mathrm{momentum} + \text{artificial viscosity}. \tag{2.12}$$

Here Dalton's law gives $p = p_\mathrm{e} + p_\mathrm{i} = 2nk_B T$, where p_e and p_i are the partial electron and ion pressures as in equation 2.9. This assumes the electron and ion temperatures equal, i.e. $T_\mathrm{e} = T_\mathrm{i} = T$. The electron contribution for the momentum equation is only taken into account in this static pressure term as the electron mass is negligible compared to the ion mass, i.e. $m_\mathrm{e} \ll m_\mathrm{i}$ and hence of no consequence for the mass-dependent terms.

While a discussion of the numerical implementation is not included in this thesis—one is instead referred to the DIV1D (Westerhof, 2019) and SD1D manuals (Dudson, 2016)—one additional comment should be made with respect to equation 2.12. In the numerical implementation of this equation an additional term has been added on the right-hand side: the artificial viscosity. As the name suggests, it is not a physical term. In following the SD1D code, it was "implemented as a diffusion of momentum in index space" (Dudson, 2016) in DIV1D to stabilise the solution. A form of artificial viscosity was first put forward by Von Neumann and Richtmyer (1950) in the context of hydrodynamic shocks. For a more detailed description of artificial viscosity and a comparison of it to alternative methods like slope limiting and spectral filters one is referred to e.g. Michoski *et al* (2016) and references therein. To get a sense of the effect that this artificial viscosity has on the DIV1D solutions one is instead directed to results section 3.2, and the discussion thereof in section 6.1.

The DIV1D energy balance can be derived from equations 2.8–2.10 by assuming that electrons and ions are in thermal equilibrium, i.e. $T_\mathrm{e} = T_\mathrm{i} = T$. Writing down the energy balances for ions and electrons

and again applying quasi-neutrality one can sum over both energy balances to arrive, after manipulation of the third term in equation 2.8 using the chain rule, at

$$\left[\frac{\partial}{\partial t}\left(\frac{3}{2}p\right)+\frac{\partial}{\partial x}\left(\frac{5}{2}pv_{\parallel}\right)-v_{\parallel}\frac{\partial p}{\partial x}\right]-\frac{\partial}{\partial x}\left(\kappa_{\mathrm{i}}\frac{\partial T}{\partial x}+\kappa_{\mathrm{e}}\frac{\partial T}{\partial x}\right)=Q_{\mathrm{i}}+Q_{\mathrm{e}}. \tag{2.13}$$

Combined with the equation of state 2.9, by absorbing the ion and electron contributions to the source term, Q_{I} and Q_{e}, into one term Q, to be determined in sub section 2.3.3, the energy balance becomes

$$\frac{\partial}{\partial t}(3nk_{\mathrm{B}}T)=-\frac{\partial q_{\parallel}}{\partial x}+v_{\parallel}\frac{\partial p}{\partial x}+Q. \tag{2.14}$$

This equates the temporal development of the internal (thermal) energy of the plasma with minus the divergence of the total heat flux $q_{\parallel}$, the work done by the pressure gradient and the sum of all energy sources and sinks Q. The total heat flux is given by

$$q_{\parallel}=5nk_{\mathrm{B}}Tv_{\parallel}-(\kappa_{\mathrm{i}}+\kappa_{\mathrm{e}})\frac{\partial T}{\partial x} \tag{2.15}$$

where we choose the thermal heat conductivities for ions and electrons, κ_i and κ_e, as the Spitzer conductivity (Cohen, Spitzer & Routly, 1950; Spitzer & Härm, 1953)

$$\kappa_{\mathrm{i,e}}=\kappa_{\mathrm{i0,e0}}T_{\mathrm{i,e}}^{5/2}\times(k_{\mathrm{B}}/e)^{7/2} \tag{2.16}$$

where $\kappa_{\mathrm{i0}}\approx 60\ \mathrm{W\,m^{-1}\,K^{-7/2}}$, $\kappa_{\mathrm{e0}}\approx 2000\ \mathrm{W\,m^{-1}\,K^{-7/2}}$ (Stangeby, 2000) and where the factor $(k_{\mathrm{B}}/e)^{7/2}$ corrects for this derivation's choice of K rather than eV, the latter of which is employed by Stangeby. Note the strong temperature dependence of the thermal conductivity in equation 2.16.

Finally, it should be noted that different from the DIV1D manual (Westerhof, 2019) and the actual implementation in the DIV1D code, this section has kept to the SI formulation of the plasma fluid equations, which means that temperatures are always in kelvin. Adapting equations 2.11–2.16 to the electron volt formulation, i.e.

$$k_{\mathrm{B}}T[\mathrm{K}]=eT[\mathrm{eV}], \tag{2.17}$$

simply means equating the Boltzmann constant k_{B} with the elementary charge e whenever it presents itself in the equations (but not *vice versa*), then using the temperature in electron volt. The results section uses eV instead of K for temperatures, because it allows for more tractable numbers. They can be readily converted to SI through a multiplication factor of $e/k_{\mathrm{B}}=11.6\times10^{3}\ \mathrm{K/eV}=1\times10^{4}\ \mathrm{K/eV}$.

2.3.2. Neutral transport, sources and sinks

This section describes the neutral particle transport and the neutral sources and sinks in DIV1D. Unlike plasma transport, where the constituent particles are stuck to the magnetic field lines and which thus lends itself well to a one-dimensional approach, neutral particles are not bound to the magnetic field lines and hence constitute an inherently two-dimensional process. Self-consistent modelling of neutral

transport requires the use of complex three-dimensional codes such as EIRENE (Reiter, 2019), which is infeasible, and moreover, unwanted for DIV1D.

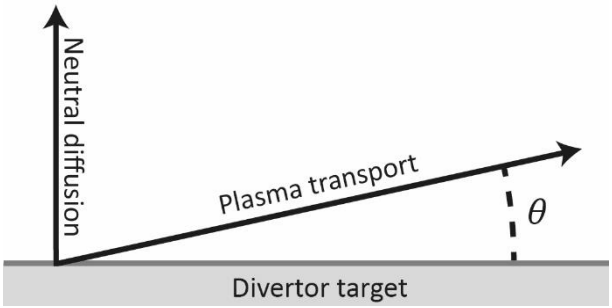

Figure 2.1. Schematic of how the directions of plasma transport and of neutral diffusion differ in a slanted-divertor configuration in the direction along the flux tub, i.e. a configuration in which the divertor sits at an angle θ with respect to the local magnetic field. This increases the neutral diffusion coefficient as seen from a magnetic–field-line frame of reference by a factor of $\sin^2\theta$. The figure is based on Nakazawa *et al* (2000).

Since it is impossible to model neutrals correctly, then, their evolution might as well be kept as simple as possible. Thermal neutrals are very cold with respect to the rest of the plasma, and hence carry but a small fraction of the system's total energy and momentum. They are consequently set to zero in DIV1D, saving the code from having to calculate the neutral momentum and energy balances. The number of neutrals, however, is of the same order as the number of plasma particles, at least near the target, so the neutral particle balance cannot likewise be ignored. The neutral particle balance in DIV1D is given as a simple diffusion equation

$$\frac{\partial n_{\mathrm{neutral}}}{\partial t} = \frac{\partial}{\partial x} D \frac{\partial}{\partial x} n_{\mathrm{neutral}} + S_{\mathrm{density,neutral}} \tag{2.18}$$

with n_{neutral} the neutral particle density, $S_{\mathrm{density,neutral}}$ the sum of all sinks and sources for the neutrals and D the diffusion coefficient as a function of plasma temperature and plasma density following (Nakazawa, 2000)

$$D = \frac{k_{\mathrm{B}} T}{m \nu_{\mathrm{cx}} \sin^2\theta}. \tag{2.19}$$

Here $\nu_{\mathrm{cx}} = n\langle\sigma_{\mathrm{cx}}\rangle$ is the charge exchange collision frequency with $\langle\sigma_{\mathrm{cx}}\rangle$ the plasma-temperature–dependent charge-exchange reaction rate, i.e. the collisional cross section integrated over the velocity spectrum, and θ the pitch of the magnetic field with respect to the divertor plane at the divertor. Why this latter factor needs to be taken into account becomes evident from figure 2.1. It is assumed that neutral diffusion happens perpendicularly to the divertor plane, which means that from the point of view of a plasma neutral particles seem to be diffusing faster than they really are. The source term for the neutrals is given as

$$S_{\mathrm{density,neutral}} = -S_{\mathrm{density}} + S_{\mathrm{puff}} - \frac{n_{\mathrm{neutral}}}{\tau_{\mathrm{neutral}}} \tag{2.20}$$

where $-S_{\mathrm{density}}$, to be discussed in more detail in the next sub section, reflects that what is a source for the plasma is a sink for the neutrals, and vice versa. S_{puff} allows for the simulation of fuelling through

a constant neutral gas puff. Finally, τ_{neutral} is a finite residence timescale, so that $n_{\mathrm{neutral}}/\tau_{\mathrm{neutral}}$ reflects a constant loss of particles given a fixed number of neutrals per unit volume.

2.3.3. Plasma sources and sinks

In section 2.2.1. the plasma fluid equations were derived. A full closure of those equations requires a description of the source terms S_{density}, S_{momentum} and Q. They are discussed in this sub section. The plasma density, firstly, is affected in two ways. Electrons and ions are removed through recombination, and replenished through ionisation, so the density source term is given as

$$S_{\mathrm{density}} = -n^2\langle\sigma_{\mathrm{rec}}v\rangle + nn_{\mathrm{neutral}}\langle\sigma_{\mathrm{ion}}v\rangle \tag{2.21}$$

where $\langle\sigma_{\mathrm{rec}}v\rangle$ and $\langle\sigma_{\mathrm{ion}}v\rangle$ are the plasma-temperature–dependent effective reaction rates for recombination and ionisation. The plasma momentum , secondly, is affected by charge exchange and, again, by recombination, both of which constitute loss channels. Since the neutral momentum has been set to zero, no momentum sources exist. The momentum source term is thus given by

$$S_{\mathrm{momentum}} = -mnv_{\parallel}(n_{\mathrm{neutral}}\langle\sigma_{\mathrm{cx}}v\rangle + n\langle\sigma_{\mathrm{rec}}v\rangle). \tag{2.22}$$

The plasma energy, finally, also only contains loss channels, since neutral energy has been likewise set to zero. The loss channels are due to charge exchange, recombination, excitation and impurity-radiative cooling, and summed to yield

$$Q = -\frac{3}{2}nk_{\mathrm{B}}T(n_{\mathrm{neutral}}\langle\sigma_{\mathrm{cx}}v\rangle + 2n\langle\sigma_{\mathrm{rec}}v\rangle) - n\big(n_{\mathrm{neutral}}L_{\mathrm{exc}}(T) + n_Z L_Z(T)\big) \tag{2.23}$$

where $n_Z = n\xi_Z$ is the impurity density with ξ_Z the impurity concentration, and $L_{\mathrm{exc}}(T)$ (which also takes into account the energy effect due to ionisation) and $L_Z(T)$ are cooling rates due to temperature-dependent electron–neutral excitation and impurity radiation respectively.

2.4. Analytical solution: two-point model

The two-point model is a marvellous starting point to understanding how the divertor couples to the core plasma via the scrape-off layer. In rather broad strokes, it lays out a typical divertor plasma scenario from which three relations can be derived that link key plasma parameters on the upstream (core) side of the scrape-off layer with those on the divertor target (downstream) side. This model, although extremely simple, readily makes insightful why we will be so focussed on 'just a few' specific plasma parameters further on in this work. Which parameters these are, and how they scale with respect to the input parameters are the subject of this section. This section closely follows Stangeby (2000).

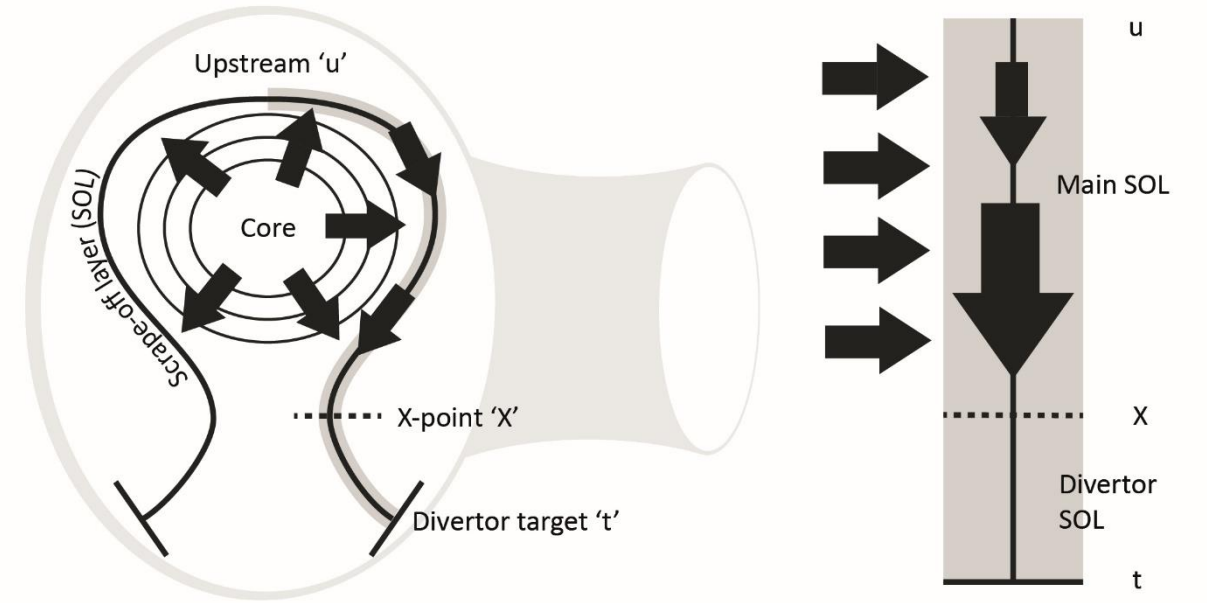

Figure 2.2. This is a schematic overview of the areas that are of key interest to understanding a fusion heat exhaust. Black lines denote the poloidal component of the magnetic field. Heat is generated in the core plasma. It then moves out purely radially (since no gradients can exist in the plane of the flux surfaces) until it hits the last closed flux surface (LCFS). Beyond this point, magnetic field lines cease to wrap around the torus indefinitely. Rather, after a finite number of toroidal turns they terminate on the divertor target. This part of the plasma is called the scrape-off layer. As parallel transport is much faster than perpendicular transport, heat now starts moving poloidally rather than radially and eventually hits the divertor target.

Both in the two-point model (2PM) as in 1D modelling, things are considerably simplified by straightening out the scrape-off layer (SOL). In this case, the heat flux builds up as the SOL is traversed from the upstream side to the target side. In case of the 2PM, the actual position of the upstream side is never really specified—it enters indirectly though, through the connection length—while in DIV1D the upstream side is chosen to coincide with the X-point position. This is motivated by the fact that the main SOL acts simply as a collecting bin to the heat flux that comes out of the core plasma. Once all heat has been collected near the X-point the entirety of it—all the radially exuded heat flux integrated over the main SOL length—enters the divertor SOL. For a given core performance this value is fixed, so it might as well be given as a boundary condition. This figure is based on figure 5.7 in Stangeby (2000).

Box 2.2. Simplifying assumptions to the two-point model

With respect to the particle balance (equation 2.6) it is assumed that

I. plasma particles cannot leave the domain perpendicularly to the flux tube;
II. there exists a thin ionisation layer directly in front of the divertor target where all wall-recycled neutrals are ionised;
III. volume recombination is negligible;
IV. plasma flow is negligible except in the thin ionisation layer directly in front of the target, where it grows from zero to match the sheath entrance speed at the end.

With respect to the momentum balance (equation 2.7) it is further assumed that

V. viscosity and friction are negligible;
VI. there exists no electric field parallel to the flux tube.

Finally, with respect to the power balance (equation 2.8) it is assumed that

VII. heat transport is conduction-dominated because of assumption IV;
VIII. the conductivity of the plasma can be described by the Spitzer conductivity;
IX. power sources and sinks are negligible.

2.4.1. Basic two-point model

Based on a set of assumptions given in box 2.2. and applying those to the balance equations 2.6–2.8 the following two-point model (2PM) can be derived as

$$2n_{\mathrm{t}}T_{\mathrm{t}} = n_{\mathrm{u}}T_{\mathrm{u}};$$

(2.24)

$$T_{\mathrm{u}}^{7/2} = T_{\mathrm{t}}^{7/2} + \frac{7}{2}q_{\parallel}\frac{L}{\kappa_{\mathrm{e0}}};$$

(2.25)

$$q_{\parallel} = q_{\mathrm{t}} = \gamma n_{\mathrm{t}}k_B T_{\mathrm{t}}c_{\mathrm{st}}.$$

(2.26)

All symbols denote the same quantities as before. Newcomers are L to denote the distance between the upstream and target positions along the magnetic field, κ_{e0} to denote the Spitzer thermal conductivity for electrons at a temperature of e/k_{B} [K] (or 1 eV), γ the sheath heat transmission coefficient and c_{st} the sound speed $\sqrt{2k_{\mathrm{B}}T/m_{\mathrm{i}}}$. The subscripts u and t refer to quantities evaluated at the upstream or target sides, respectively, while the subscript $\parallel$ somewhat superfluously emphasises that the heat flux is parallel to the domain.

Three equations 2.24–2.26 allow for three unknowns: T_{t}, T_{u} and n_{t}. Meanwhile, $q_{\parallel}$, n_{u} and L are machine parameters, with the first two determined by core performance and the last one an engineering parameter. Finally, γ and κ_{e0} are constants. It is interesting to know how the three unknowns depend on our machine parameters. Firstly, since equation 2.25 scales so strongly with the temperature and noting that the temperature drops considerably over the course of the scrape-off layer ($T_{\mathrm{u}} > T_{\mathrm{t}}$), equation 2.25 can be simplified to give

$$T_{\mathrm{u}} = \left(\frac{7}{2}\frac{q_{\parallel}L}{\kappa_{\mathrm{e0}}}\right)^{2/7}.$$

(2.27)

This equation shows that T_{u} is only very weakly dependent on the heat flux, or any other parameter for that matter. As a numerical example, a factor of ten decrease in the heat flux not so much as halves the upstream temperature. This means that scaling up the power production of a fusion device by some two orders of magnitude—which is in the ballpark of the projected jump from the Joint European Torus (JET), today's flagship experimental fusion device, to the International Thermonuclear Experimental Reactor (ITER), currently under constructing in France—will only increase the upstream temperature by a factor of four[3].

Meanwhile, equations 2.24, 2.26 and 2.27 can be combined to show that the target temperature, the second unknown, scales as

$$T_{\mathrm{t}} = \frac{2m_{\mathrm{i}}}{\gamma^2 e^3}\frac{q_{\parallel}^2}{n_{\mathrm{u}}^2}\left(\frac{7}{2}\frac{q_{\parallel}L}{\kappa_{\mathrm{e0}}}\right)^{-4/7} \propto q_{\parallel}^{10/7}L^{-4/7}n_{\mathrm{u}}^{-2}.$$

(2.28)

[3] One could argue that this ignores the increase in the connection length, so that the temperature jump would be bigger in reality. However, a growing machine size also increases the scrape-off layer's footprint (which grows as the major radius) in turn driving down the heat flux, while the connection length for a given pitch angle will also grow as the major radius. Both effects cancel out.

T_t is one parameter that we would ideally minimise, as was discussed in the previous chapter. However, T_t scales rather unfavourably with $q_\parallel$. This is a pity because for a given machine size the total power produced in the core is linearly proportional to $q_\parallel$, so from the point of view of maximising power production we find that more $q_\parallel$ is better. We would much prefer not to turn that particular knob to drive down T_t. Can one of the remaining parameters compensate? Putting too much stock in the connection length L is ill-advised for several reasons. For one, T_t scales rather poorly with growing L. For another, a big part of L is determined at device assembly and can thus not readily be used to accommodate plasma fluctuations with typical timescales of less than forty years. This really only leaves one option: the upstream density n_u. T_t drops off rather quickly as a function of it, and since the plasma density relates to the fuelling rate, it is also a parameter that can be influenced during discharges.

Equations 2.24, 2.27 and 2.28 can be combined to yield the scaling for the third unknown

$$n_t = \frac{\gamma e^2}{2m_i}\frac{n_u^3}{q_\parallel^2}\left(\frac{7}{2}\frac{q_\parallel L}{\kappa_{e0}}\right)^{6/7} \propto q_\parallel^{-8/7} L^{6/7} n_u^3, \tag{2.29}$$

which shows that n_t scales particularly strongly with n_u. Finally, there is one derived quantity that begs a closer look, because it is a key parameter for target erosion. This is the particle flux Γ_t onto the target, which is found by dividing equation 2.26 by $\gamma k_B T$ and inserting 2.28 to yield

$$\Gamma_t = \frac{\gamma e^2}{2m_i}\frac{n_u^2}{q_\parallel}\left(\frac{7}{2}\frac{q_\parallel L}{\kappa_{e0}}\right)^{4/7} \propto n_u^2 q_\parallel^{-3/7} L^{4/7}. \tag{2.30}$$

The main takeaway here is that n_u, which is the prime candidate for driving down T_t according to equation 2.28, has an unfavourable effect on the particle flux onto the target. Since $\Gamma_{sput} = Y\Gamma_t$ it will then depend on the nature of the sputtering process, embodied by the sputtering yield Y, whether this is a problem.

2.4.2. Extending the two-point model towards detachment

As was described in section 1.4, two requirements for detachment to occur are a loss of power and a loss of momentum (pressure) between the upstream and divertor sides of the tokamak. These are minimal requirements for detachment to occur; if either of the two is missing, no detachment can occur (Dudson *et al*, 2019). As a result, a discussion of detachment cannot take place within the framework of the basic two-point model, which does not allow for momentum or power losses by design. Two-point models come in different levels of complexity, but for the understanding of detachment a most simple version suffices. Power loss is taken into account in equation 2.26 by having only a part—one minus the power loss fraction f_{power}—of the upstream heat flux $q_\parallel$ reach the target,

$$(1 - f_{power})q_\parallel = q_t = \gamma k T_t n_t c_{st}. \tag{2.31}$$

Likewise, momentum loss can be inserted into equation 2.24 as

$$2n_t T_t = (1 - f_{mom}) n_u T_u, \tag{2.32}$$

where f_{mom} denotes the fraction of momentum that is lost between upstream and target[4]. Finally, the effect of convection is taken into account. As an effect that tends to reduce gradients temperature, according to Stangeby, it acts on the conductive term, which scales with ∇T. Introducing a loss fraction due to convection, f_{conv}, equation 2.25 becomes

$$T_{\mathrm{u}}^{7/2} = T_{\mathrm{t}}^{7/2} + \frac{7}{2}(1 - f_{\mathrm{conv}}) q_{\parallel} \frac{L}{\kappa_{\mathrm{e0}}}. \tag{2.33}$$

Retracing the steps of the previous section, multiplication factors can be determined for equations 2.28–2.30 that correct for the loss effects in equations 2.31–2.33, which yields

$$T_{\mathrm{u}} \propto (1 - f_{\mathrm{conv}})^{2/7}; \tag{2.34}$$

$$T_{\mathrm{t}} \propto \frac{\left(1 - f_{\mathrm{power}}\right)^2}{(1 - f_{\mathrm{mom}})^2 (1 - f_{\mathrm{conv}})^{4/7}}; \tag{2.35}$$

$$n_{\mathrm{t}} \propto \frac{(1 - f_{\mathrm{mom}})^3 (1 - f_{\mathrm{conv}})^{6/7}}{\left(1 - f_{\mathrm{power}}\right)^2}. \tag{2.36}$$

For given values of n_{u}, $q_{\parallel}$ and L, the dependent parameters T_{t}, T_{u} and n_{t} only depend on these loss fractions. Just as for all other parameters, the upstream temperature appears mostly impervious to the introduced loss term, depending on the convective fraction only slightly. The target temperature, meanwhile, depends strongly on both power and momentum losses. An increase in the former gives rise to as strong decrease of the target temperature while an increase in the latter gives rise to a strong increase of the target temperature. The target density, not unintuitively, shows an inverse tendency with respect to the target temperature. Now it is the power losses that cause an increase and the momentum losses that cause a decrease. Note that the momentum effect is even stronger than for the temperature case. Although the convected fraction is never the dominant effect it does influence both target temperature and density. It slightly increases the former—congruent with the assumed tendency of the convective component to reduce the temperature gradient (as it also lowers the upstream temperature)—while it simultaneously decreases the latter, albeit at a somewhat faster rate. Finally, the particle flux scales with

$$\Gamma_{\mathrm{t}} \propto \frac{(1 - f_{\mathrm{mom}})^2 (1 - f_{\mathrm{conv}})^{4/7}}{1 - f_{\mathrm{power}}}. \tag{2.37}$$

An interesting result of this is that power loss of itself is possibly detrimental to the rate of erosion, depending on the temperature-dependence of the sputtering yield.

[4] Those familiar with the original derivation by Stangeby (2000) may observe here that f_{mom} is chosen to reflect the momentum loss fraction rather than the fraction that does reach the target. This was chosen as such to come to formulations similar to that for the power loss. For a similar reason, I use $(1 - f_{\mathrm{conv}})$ with f_{conv} the fraction lost to convection, rather than the conducted fraction f_{cond} itself.

3. Code verification

Over the course of doing this thesis work there were occasions on which we encountered (at the time) difficult to explain behaviour. It was not always immediately clear whether these behaviours were a result of the physical model itself or of some numerical fluke. One feature of the DIV1D solutions that jumps to mind, here, is the way in which the solutions always seemed to decay in a similar damped oscillatory fashion after an initial perturbation of the heat flux. Peak ahead at figures 3.8 and 3.9 for a second to see what this means in practice. First observed in the context of ELM[5] simulations, discussed in chapter 5, and initially treated as a result of those ELMs, these damped oscillations later turned out to be much more universal to DIV1D simulations.

To lend strength to these observations, if not to fully preclude the numerics from having a hand in them, some of the more obvious numerical effects were tested. In section 3.1 the results of a grid convergence scan are shown and discussed, which was performed for a weakly detached plasma—with a target temperature between $0.1\ \mathrm{eV}$ (which is the lowest value that DIV1D allows) and $10\ \mathrm{eV}$—in a range between 100 and 3500 cells. It shows the behaviour of the free 2PM parameters—upstream temperature, target temperature and target density—as well as of the momentum and power loss fractions from the extended 2PM. In section 3.2 the results of two artificial viscosity scans are shown and discussed. The first was executed for the same weakly detached plasma as above, and the second for a more deeply detached plasma in which the detachment front has receded several meters from the target. These two cases—the $n_\mathrm{u} = 3.9 \times 10^{19}\ \mathrm{m}^{-3}$ weakly detached scenario and the $n_\mathrm{u} = 5.0 \times 10^{19}\ \mathrm{m}^{-3}$ deeply detached scenario— are the starting points of several other simulations presented in chapters 4 and 5.

3.1. Grid convergence

Most of the DIV1D results shown in this thesis were generated using either 500-cell or 1000-cell grid, that is non-uniformly distributed. The fineness of the grid increases towards the target. Calculations were capped at this degree of discretisation to allow for sufficiently fast calculations. That is, calculations with run times on the order of a few hours to a few days. Ideally, this 1000-cell spatial resolution should be in a regime where the solutions have become independent from the spatial resolution, i.e. a regime in which further increasing the number of grid points does not further affect the solution. To determine whether this was the case, a physics scenario with identical input parameters ($n_\mathrm{u} = 3.9 \times 10^{19}\ \mathrm{m}^{-3}, q_\parallel = 50\ \mathrm{MW\ m}^{-2}$ and $L = 20\ \mathrm{m}$ under full wall recycling—which are similar to the parameters used in Dudson *et al* (2019)) was tested over a range of resolutions.

Figures 3.1 through 3.4 show time-averaged upstream and target parameters from the 2PM along with the power and momentum loss fractions that are important for the extended version. They are time-averaged over the last $2\ \mathrm{ms}$ of the simulation. The error bars give the standard deviation of this averaging operation, which can also be read as the root mean square (RMS) of the remaining decaying oscillations. They are next to negligible when compared to the trends that are visible in the data.

[5] ELMs are Edge-Localised Modes. They are type of instability of the core plasma that dumps a significant fraction of heat and particles into the SOL over the course of milliseconds. The ensuing heat flux to the divertor target plates is huge and poses a significant risk to the survival of the target. They'll be discussed in more detail in chapter 5, which elaborates on their addition to DIV1D.

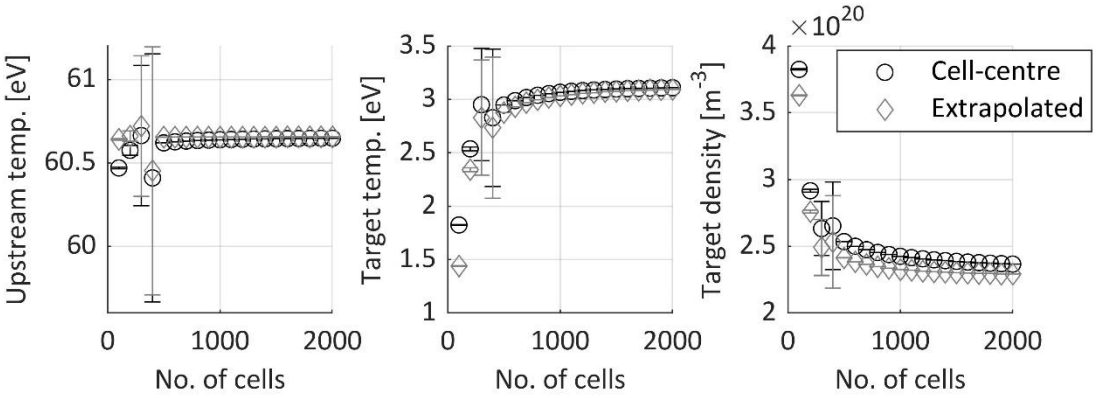

Figure 3.1. The independent parameters of 2PM as functions of the spatial resolution. Parameters were averaged over the last two milliseconds of otherwise identical ten-millisecond DIV1D simulations with the error bars denoting the standard deviations to those averages. Black circles denote parameters evaluated at the first or final cell, while grey diamonds denote extrapolations to the left (0-m) and right (20-m) boundaries of the domain based respectively on values in the first and last five cells. All simulations were started on the (inter- or extrapolated) values obtained from the 1000-cell simulation. Two simulations, with grids of 300 and 400 cells, have significantly larger error bars than the rest. Visual inspection of the time traces indicated that these simulations experienced numerical instabilities.

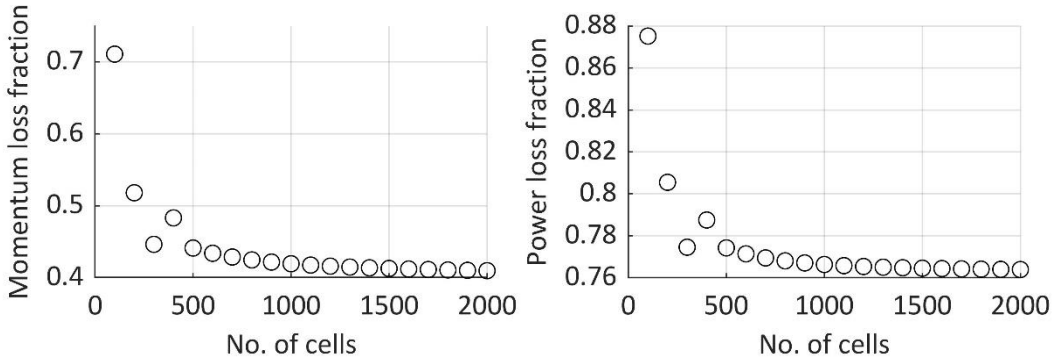

Figure 3.2. The power and momentum loss fractions as functions of the spatial resolution for the same simulations as in figure 3.1. The two values that do not sit in the trend all that nicely are the same 300-cell and 400-cell simulations that have the large error bars in the free upstream and target parameters.

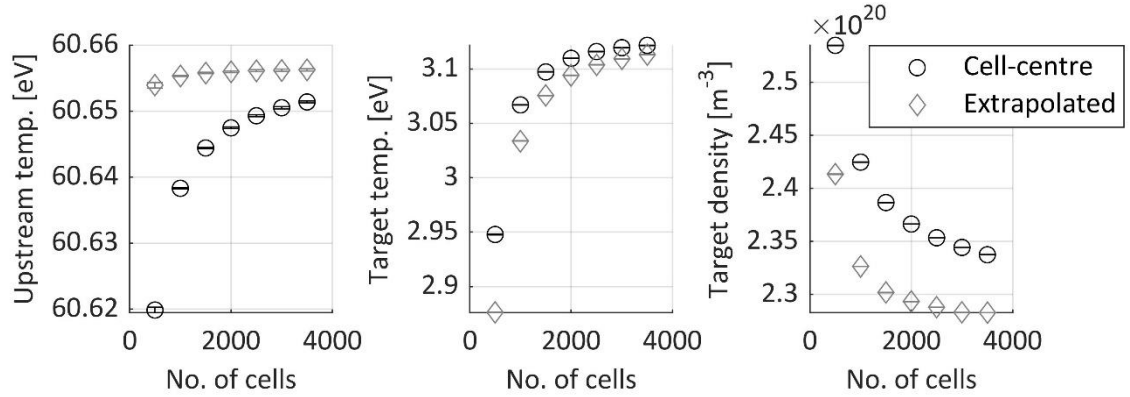

Figure 3.3. The independent parameters of 2PM as functions of the spatial resolution for even higher cell counts and taking bigger steps. Although the trend seems unabating at first, note that the vertical axes span an order of magnitude less than they did previously.

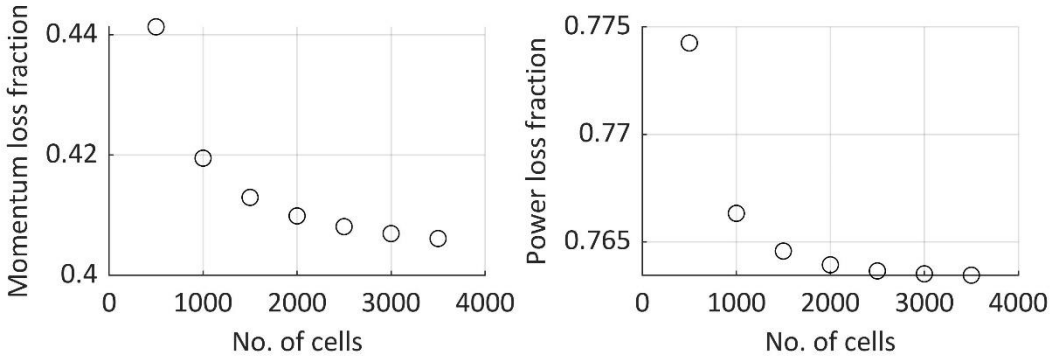

Figure 3.4. The power and momentum loss fractions as functions of the spatial resolution for the same simulations as in figure 3.3. Similar to the free parameters, the incremental changes in the momentum and power loss fraction are significantly less at higher spatial resolutions.

Figures 3.1 and 3.2 show results from a spatial resolution scan that spans the range between 100 and 2000 cells in 100-cell increments for a non-uniformly distributed grid which samples with more precision near the target than near the upstream. Two metrics are plotted. The values in the final (for target parameters) or first cells (for upstream parameters) are given in black circles. The positions along the flux tube of these first and final cells change as a function of grid spacing, however, and they thus introduce an extra dependence on the resolution. Although effects are expected to be slim—parameters are rather constant for the first few meters on the upstream side while they are sampled very finely close to the target—parameters from the first or last five grid cells were additionally extrapolated to the beginning ($x = 0$) and end ($x = L$) of the domain. For all but the coarsest grids (≤ 300 cells) these five points were located between the density peak and the target. These extrapolated values are denoted by grey diamonds. The effect of this extrapolation is particularly noticeable for the upstream temperature (although the absolute differences are small in either case). While a small dependence on the resolution is observed in the upstream temperature—tending towards higher values for higher resolutions—in the single-cell scenario, this effect all but disappears in the extrapolated scenario. This is more readily visible in figure 3.3, which depicts a scan at 500-cell increments in the range between 500 and 3500 cells for the same physics scenario as in figure 3.2 (with the first four data points in the figure 3.3 equal to those in figure 3.2).

In the case of the target parameters a trend remains towards higher temperatures and lower densities. This trend does not abate when moving from the single-cell to the extrapolated perspective. Similar trends are observed for the momentum and power loss fractions, which were calculated as the losses integrated over the flux tube and divided by the upstream parameters and for which it holds that $0 \leq f_{\mathrm{mom}} \leq 1$ and $0 \leq f_{\mathrm{pow}} \leq 1$. The momentum loss fraction is more strongly influenced by varying the resolution than the power loss fraction, the former changing by 0.3 from 100 to 2000 grid points and the latter by only 0.1. The same is observed from 500 to 3500 grid points, over which the momentum loss fraction decreases by 0.04 while the power loss only decreases by 0.01. For all four parameters it is observed that the differences between increments become smaller at higher resolutions, which indicates that they should converge at some point. However, they do so less readily than the upstream temperature.

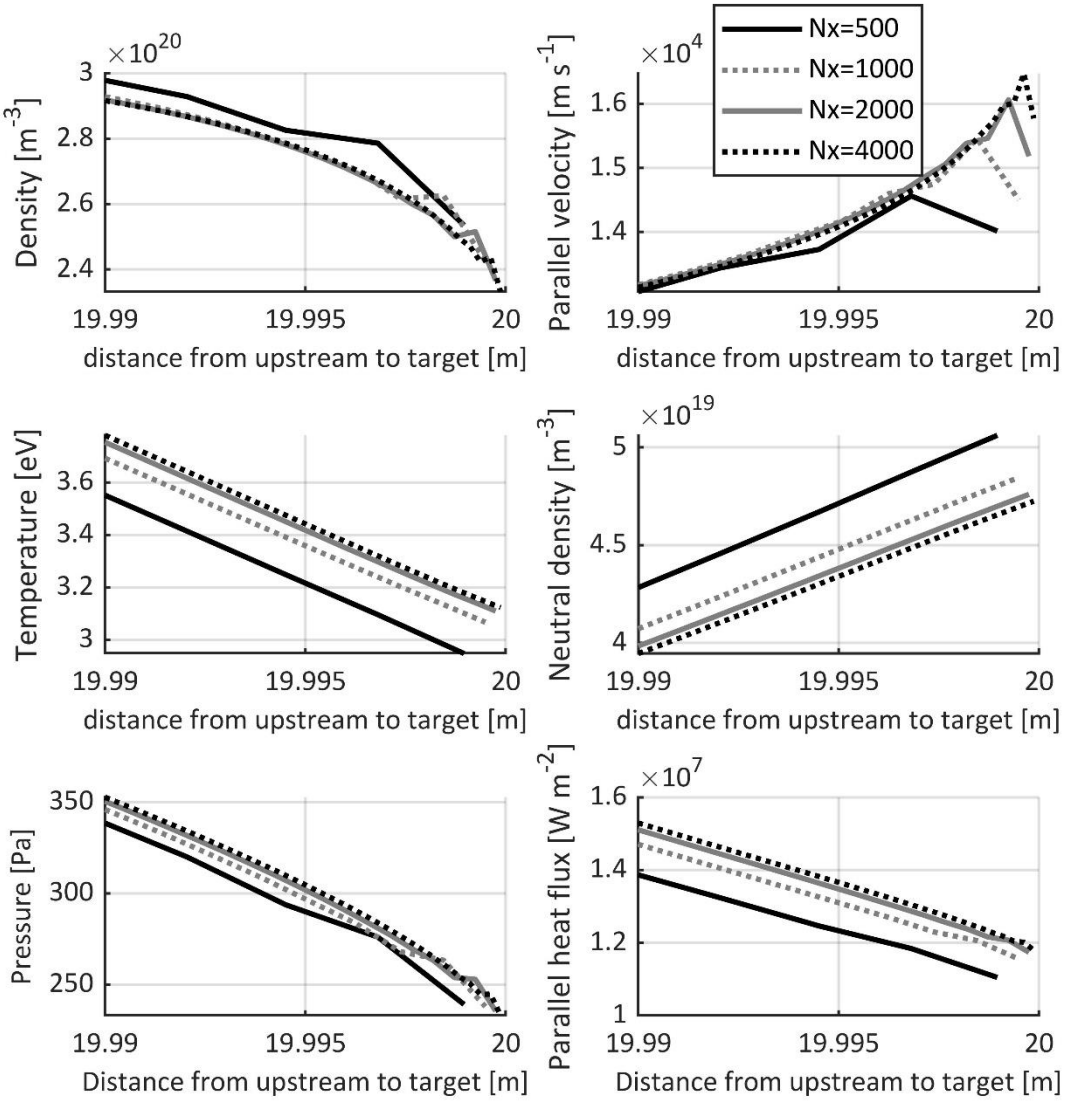

Figure 3.5. Profiles of several plasma parameters in the last centimetre before the target for different spatial resolutions. A kink is visible across all resolutions in the penultimate data point before the target in some of the parameters. The plasma density, parallel velocity, pressure and parallel heat flux are visibly affected. The temperature and neutral density do not appear to be influenced.

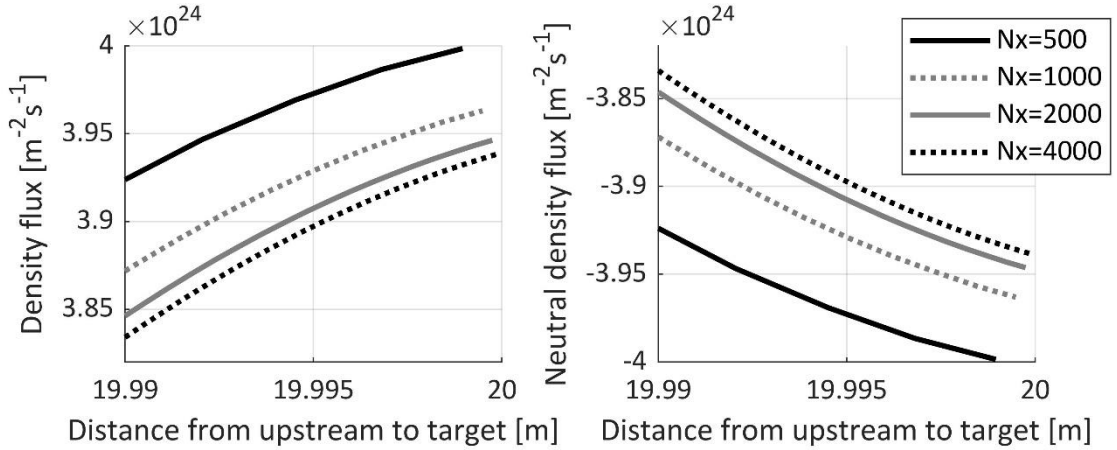

Figure 3.6. Profiles of the (plasma) density and neutral density fluxes in the last centimetre before the target for different spatial resolutions. The kink that was visible in the density and the velocity no longer appears here.

Finally, we zoomed into the last centimetre before the target to see how solutions behaved there. The parallel profiles of the density and temperature are plotted along with several others in figure 3.5. What catches the eye is a kink just before the target. It appears in the density, the parallel velocity, the pressure and the heat flux, but not in the temperature or the neutral density. Furthermore, this kink is always centred around the penultimate data point in a given simulation. Note that the previously discussed dependence of the final grid cell position on the resolution is the cause for the varying termination points of graphs.

Because this kink is always coincident with the data point prior to the last, it would appear to be a numerical rather than a physical artefact. Although it is in a really narrow strip of the domain, the changes are rather strong. Taking for example the 2000-cell simulation, the parallel velocity makes a jump of 6% of the maximally attained value across the last two data points before the target. The plasma density, meanwhile, jumps by 5% of its maximally attained value across those last two data points. This effect is no longer visible when the density flux is plotted rather than the density itself. Refer to figure 3.6.

3.2. Artificial viscosity

DIV1D is not always free of numerical instabilities. One of the possible ways to ameliorate their effect is to increase the artificial viscosity η_{art}, which could sometimes be seen to improve the stability of the solutions. As explained in section 2.3.1, it is an addition to the right-hand-side of the momentum balance "as a diffusion of momentum in index space" (Dudson, 2016). In discretised form it takes the shape

$$\frac{\partial}{\partial t}[mnv_{\parallel}(\mathrm{ix})] = \cdots + \eta_{\mathrm{art}}[v_{\parallel}(\mathrm{ix}+1) - 2v_{\parallel}(\mathrm{ix}) + v_{\parallel}(\mathrm{ix}-1)] \tag{3.1}$$

where ix is the index of a grid cell, which runs from 1 to the grid size. Because η_{art} is a decidedly non-physical parameter, however, one needs to be careful in employing this feature. Changes to the η_{art} should, ideally, not affect any physics parameters. If they do, that means that the solution depends on something that does not exist in nature, which does not speak well for the model. The η_{art}, then, should be in a range where changes to it do not significantly alter physical parameters.

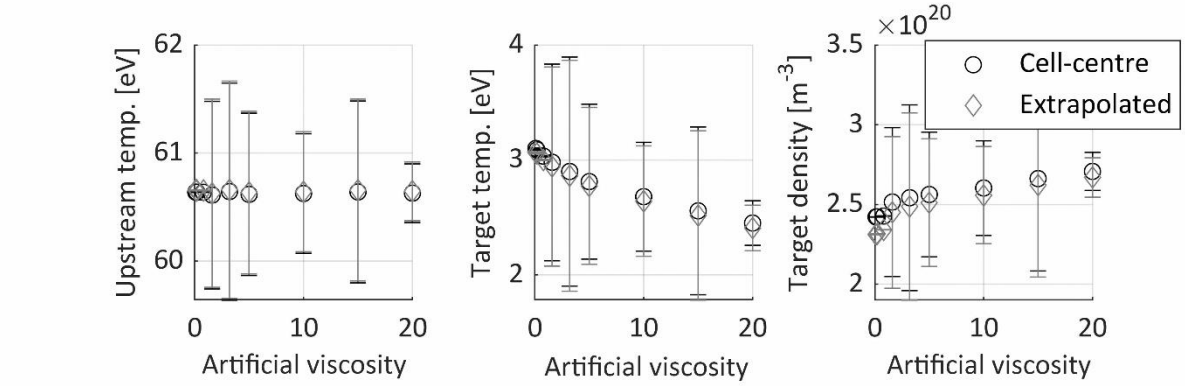

Figure 3.7. The independent parameters of the 2PM as functions of the artificial viscosity for the $3.9 \times 10^{19}\ \mathrm{m}^{-3}$ scenario. Data points represent time-averaged quantities over the full simulation duration while their error bars denote the standard deviations. This simulation duration is increased from 10 ms to 50 ms for simulations with an artificial viscosity of more than 0.8 (the third data point from the left) because simulations were not seen to converge sufficiently over 10 ms. As the error bars show, a 50-ms extension did not do much to solve this. However, time traces do seem to indicate that some sort of

quasi-steady state is reached in some simulations. Higher artificial viscosities seem to have very little effect on the mean of the upstream temperature but they do appear to affect the mean target temperature and the mean target density. The effect is opposite to that of higher grid resolution. Higher artificial viscosity causes a drop of the target temperature and an increase in the target density.

To determine the effect of the artificial viscosity on solutions, two parameters scans were performed. One was performed for the same weakly detached scenario that served as the basis for the grid convergence scan in the previous section: the $n_{\mathrm{u}} = 3.9 \times 10^{19}\ \mathrm{m}^{-3}$ scenario. The other was performed for a fully detached scenario—with the detachment front having receded several meters from the target—at an upstream density of $n_{\mathrm{u}} = 5.0 \times 10^{19}\ \mathrm{m}^{-3}$ for parameters that were otherwise identical to the weakly detached scenario. Both parameter scans were performed for the same range of values of the artificial viscosity, which spanned from 0.1–20.

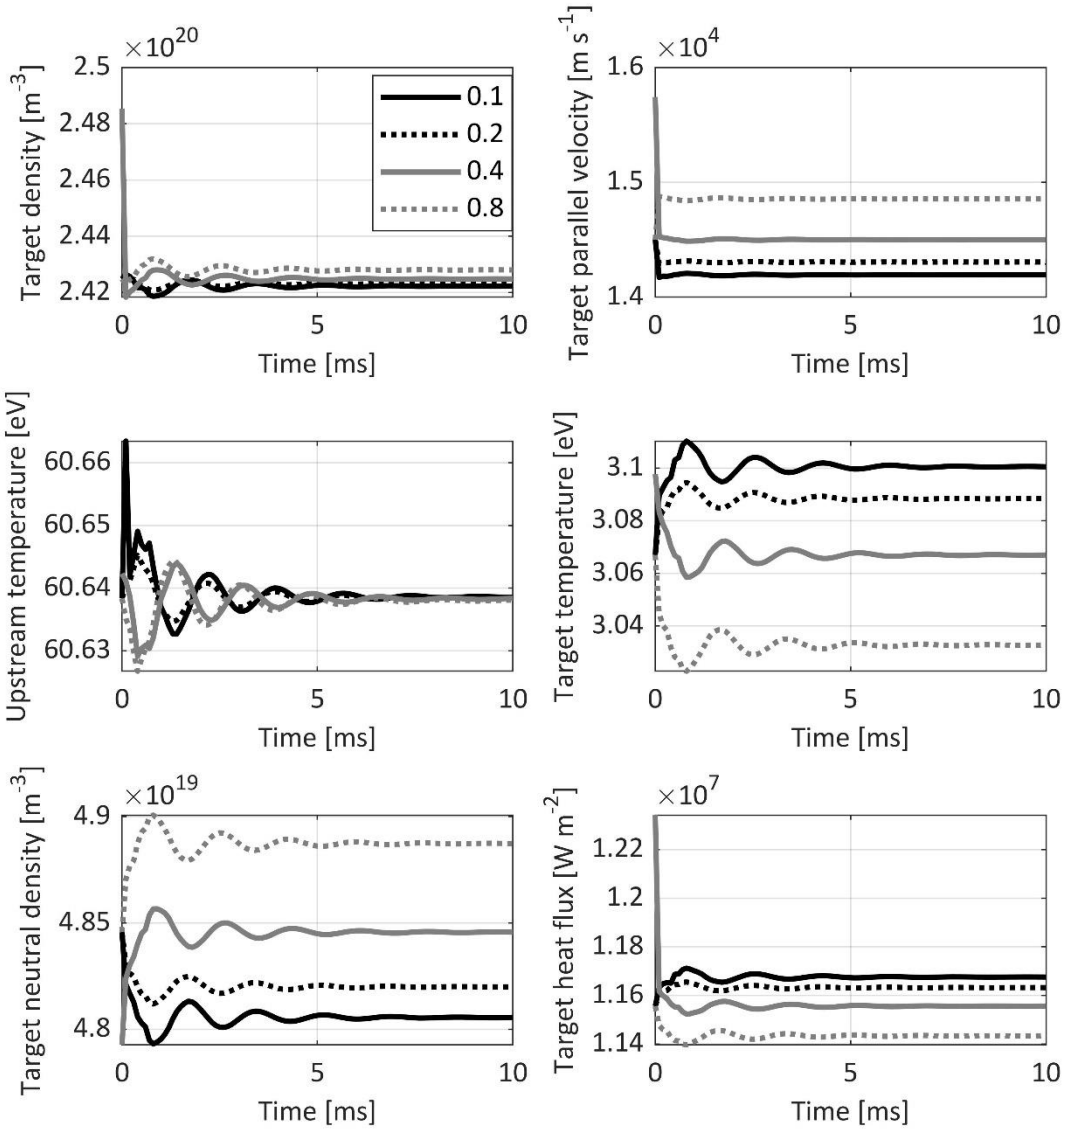

Figure 3.8. Time traces of several target parameters and the upstream temperature plotted for several values of the artificial viscosity for the $3.9 \times 10^{19}\ \mathrm{m}^{-3}$ scenario. The same is observed as in the previous figure. Higher viscosities tend to increase the target temperature while at the same time lowering the target density. The coincides with an increasing target heat flux, decreasing target plasma velocity and decreasing neutral density. As observed before, the upstream temperature is unaffected.

It was found that the increasing η_{art} does not necessarily increase the stability of the solution. For the weakly detached scenario, nicely stable results were obtained for results ranging from 0.1 to 0.8. Beyond this value, instabilities seemed to become far more present. This is clearly visible in figure 3.7, in which the error bars (indicating the standard deviation of the time averaging procedure) become much larger starting at 1.6 artificial viscosity. The standard deviation can also be interpreted as the root mean square (RMS) value of the fluctuation amplitude, and thus gives a good sense of how unstable the solution is. Even though the fluctuations are large there seems to be a trend of higher artificial viscosity relating to lower target temperatures and higher target densities. This trend is observed to hold in the range of η_{art} from 0.1 to 0.8. It somewhat difficult to see in figure 3.7 as all these values are in the left margin of the graphs. The trend is more clearly visible in figure 3.10. Lingering at figure 3.7 for one more moment, we observe that the same trend seems to propagate into the regime between 1.6 and 20. Strictly speaking, though, error bars are too big to make that observation.

Focussing on the range from $0.1-0.8$, which is the most relevant because most of the calculations in this thesis were done in that range, we observe a system that, initially out of equilibrium, converges to an equilibrium solution through a series of increasingly more damped oscillations. These oscillations are typical of a DIV1D simulation, occurring both in steady-state simulations and dynamics simulations. These oscillations are also observed directly in the wake of numerical instabilities, such as the example in figure 3.9 indicates. Although the artificial viscosity was set to an exceedingly high value of 20, these oscillations remained.

What also stands out from figure 3.8 is that although the exact same physical parameters have been inputted for all simulations, the solutions do not exactly overlap. This shows that there is at least some effect of the artificial viscosity on the solutions. This need not be a problem per se, but it does beg the question: how big is the effect? To find this out, we zoomed in on the low-viscosity part of figure 3.7. Specifically, we zoomed in on the target temperature and density, because the third free 2PM parameter, the upstream temperature, is seen to agree to within less than one hundredth of an electron volt for artificial viscosities between $0.1-0.8$ in figure 3.8.

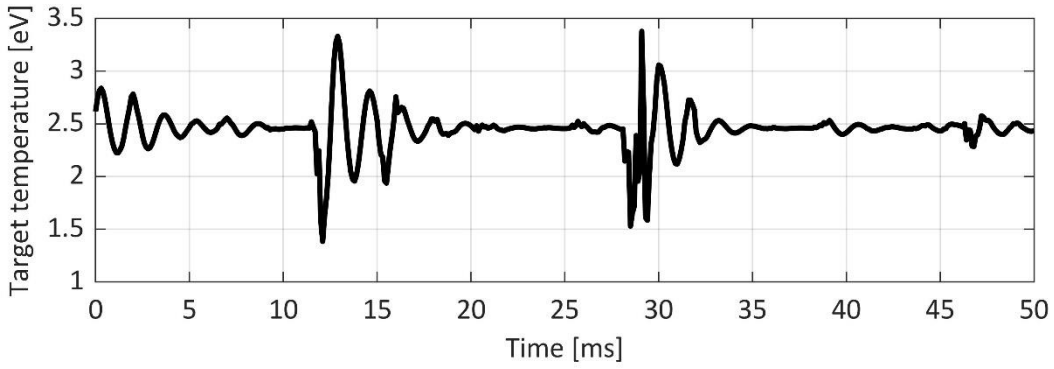

Figure 3.9. Time trace of the target temperature for the $3.9 \times 10^{19}\ \mathrm{m}^{-3}$ scenario with the artificial viscosity set to 20. The obtained target temperature is more than $0.5\ \mathrm{eV}$ lower than those obtained for artificial viscosities below 1. However, the same oscillatory behaviour is obtained regardless. And at a comparable frequency at that. Looking at the graphs, we observe around three oscillations in a 5-ms time window, both here and in the time traces provided in figure 3.8.

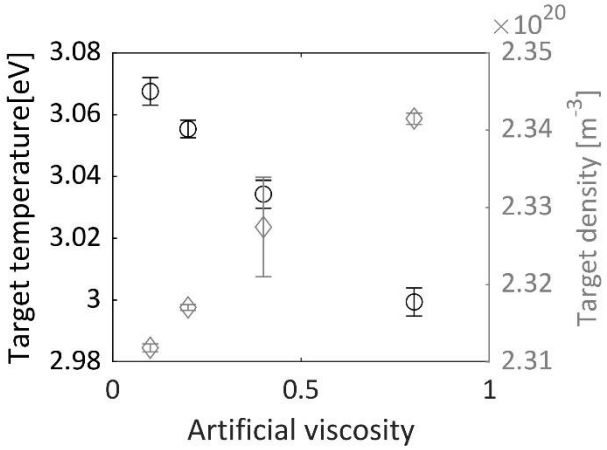

Figure 3.10. A zoom-in on the middle and right graphs in figure 3.7 that depicts the target temperature (left axis, black circles) and the target density (right axis, grey diamonds) as functions of the artificial viscosity in the range between 0 and 1. The dependence is linear, at least qualitatively speaking. Extrapolating the line to zero and one suggests that having the artificial viscosity set between zero and one changes the target temperature by a maximum of 3% with respect to the zero-viscosity case. The target density, over this same range, differs by a maximum of 2% with respect to the zero-viscosity case. The effect on the pressure is even smaller (since the temperature and density tend in opposite direction) at about 1%. The relatively large error bar for the 0.4 data point might have to do with the fact that all other measurements were started on the 0.4 converged result, whereas the 0.4 results itself was taking from the grid convergence parameter scan, with different restart conditions.

The target temperature and density were plotted against the artificial viscosity in figure 3.10. It is observed that these parameters, respectively, seem to drop and rise about linearly in this range of the artificial viscosity. Extrapolating that tendency to zero artificial viscosity, we found a target temperature of around $3.08\ \mathrm{eV}$ and a target density upwards of $2.30 \times 10^{20}\ \mathrm{m}^{-3}$. On the right side of the domain, meanwhile, values of $2.98\ \mathrm{eV}$ and $2.35 \times 10^{20}\ \mathrm{m}^{-3}$ were obtained. That means that the maximum effect of η_{art} in the range between 0 and 1 is somewhere on the order of 3% of the zero-viscosity value for the target temperature and 2% for the target density.

So what happens when this same parameter scan is executed for a deeply detached plasma? To determine the effect that the artificial viscosity has in those cases, it is no longer informative to look at any target parameters as we did for the weakly detached scenario. After all, in a deeply detached scenario the heat flux has reached zero way before the target and as such the temperature of the plasma will have dropped to numerical zero (DIV1D forces the target temperature to stay at a prefixed non-zero number, typically $0.1\ \mathrm{eV}$, to aid with numerical stability) and the density to negligible values that tell us nothing relevant about the state of the plasma further upstream[6]. As an alternative, for this parameter scan the position of the detachment front was tracked. The position of the detachment front, here, is taken to mean the position of maximum plasma density, differently from what is done in chapter 5, which tracks the position where the temperature drops below $10\ \mathrm{eV}$ for the first time—remember, this is the temperature that (at the target) indicates the onset of detachment according to Stangeby (2018). The position of maximum density was found to move in tandem with that 10-eV position, though, so changes in their respective positions are similar although their absolute positions are not the same. The maximum density is (slightly) easier to compute, since there is only maximum density while there are potentially many sub-10-eV positions along the divertor leg.

[6] To picture this, look ahead at figures 4.11, which contains several weakly detached plasmas (but also some attached ones with $T_{\mathrm{t}} > 10$ eV), and 4.12, which features a fully detached plasma in which the target heat flux has dropped to zero. See pages 55 and 57 respectively.

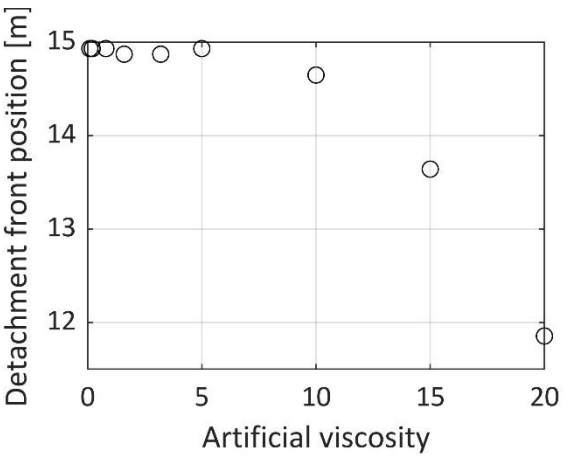

Figure 3.11. Position of the detachment front, measured as the peak plasma density, as a function of the artificial viscosity for the $5.0 \times 10^{19}\ \mathrm{m}^{-3}$ scenario. Values are almost identical up to 0.8 artificial viscosity. Between 0.8 and 5 they still do not change by a lot, but after this the effect of the artificial viscosity can no longer be ignored. The detachment front recedes by several meters for a change in artificial viscosity of 5 to 20. Note that the final result, at 20 artificial viscosity, is not yet fully converged.

Figure 3.11 shows that the position of the detachment front does not change up to an artificial viscosity of 0.8. Slight variations are observed between that point and 5, but then the drop begins in earnest: the position of the detachment front as a function of moves by at least three meters over the course of an increase in artificial viscosity of 15. The left side of figure 3.12 shows what happens to the upstream temperature for that same scan. While the upstream temperature did not appreciably change over the entire range of artificial viscosities in the weakly detached scenario, something does appear to be happening now. Starting just shy of $56\ \mathrm{eV}$ the temperature stays stable initially but then starts to drop in tandem with the shifting detachment front. Over the entire scan the upstream temperature drops by as much as 8%.

Much of this difference disappears, however, when information in figure 3.11 and the left side of figure 3.12 are combined with the 2PM scaling for the upstream temperature, $T_\mathrm{u} \propto L^{2/7}$. Combining this information is done by taking the connection length L to coincide with the position of the detachment front with respect to the upstream position L_front, i.e. the parameter plotted in figure 3.11. By dividing this parameter through a base value—L_front for the 0.1 artificial viscosity simulation in this case—a 2PM-scaled variant can be calculated as

$$T_{\mathrm{u,scaled}}(\eta_\mathrm{art}) = T_\mathrm{u} \times \left(\frac{L_\mathrm{front}(\eta_\mathrm{art})}{L_\mathrm{front}(0.1)} \right)^{-2/7}. \tag{3.2}$$

Plotting this $T_{\mathrm{u,scaled}}$ as a function of the artificial viscosity, we obtained the right graph in figure 3.12. The spread in the upstream temperature is greatly reduced, to less than $0.5\ \mathrm{eV}$. The percentage change between any two data points now stays below 1% which is much lower than the 8% that was observed for T_u itself. This would bring observations better in line with the weakly detached scenario.

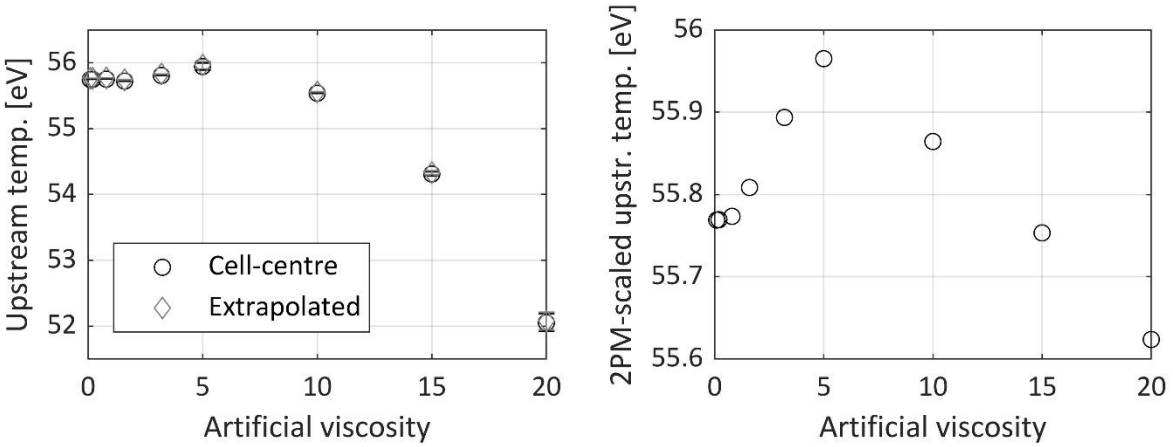

Figure 3.12. (Left) Upstream temperature as a function of the artificial viscosity for the $5.0 \times 10^{19}\ \mathrm{m}^{-3}$ scenario. Where this parameter would not budge in the $3.9 \times 10^{19}\ \mathrm{m}^{-3}$ it does change in case of the deeply detached scenario, dropping by as much as 8%. **(Right)** A combination of the data from figures 3.11 and the left-hand graph in this figure, this graph shows the effect of scaling the upstream temperature with the position of the detachment front according to equation 2.27. By doing so, the spread between the data points drops from 8% to less than 1%. Note, as in figure 3.10, the last data points on both left and right are of an unconverged run.

4. Radial power losses

This chapter discusses the results obtained to answer the question:

> What physics mechanisms drive the differences between DIV1D and SOLPS-ITER, and can they be accounted for in 1D?

Although the answer is already given away in the chapter title, this chapter shows the considerations that lead up to the implementation of radial power losses in DIV1D, as well as their eventual effects on the solutions.

To answer this question required several key ingredients. At the very least, it required one to understand the reference system: DIV1D. This requirement has in part been met by section 2.3 of the previous chapter, which detailed the physics basis of DIV1D. Taking DIV1D as the anchor point, two methods of comparison were used with varying levels of success. The idea at the onset was to compare DIV1D and SOLPS-ITER on a per-source, per-sink and per-flux basis and in that way answer the question of what the dominant physics mechanisms were that drove the difference. This method of comparison sounded rather straightforward at the beginning, and showed promising progress initially. However, over the course of time it turned out that SOLPS-ITER looks a lot more straightforward to the untrained than to the semi-trained eye. Adequately understanding the inner machinations of SOLPS-ITER as a sub question to one of two research questions proved beyond the scope of this project. Concretely, the initial attempt to reconstruct the full SOLPS-ITER power balance based on all underlying sources, sinks and flux faltered.

That realisation resulted in the second method of comparison. Rather than comparing either system on a first-principles basis, we decided to take a more phenomenological tack. After all, the intention was never to understand the physics differences just for the sake of understanding the physics differences. Ultimately, those differences were only relevant insofar as they could be used to correct DIV1D to yield better matches between either code's output. That was a good choice because it then turned out that the information extracted from SOLPS-ITER was sufficient. This chapter presents those findings as well as the way that they were added to DIV1D.

Section 4.1 discusses a SOLPS-ITER simulation of TCV that was used as a reference for comparison: shot #150662 (Wensing, 2020). Similar simulations are described in Wensing *et al* (2020). Section 4.1 introduces the parameter profiles to key physical parameters and it describes the process to obtain DIV1D-compatible input parameters from the SOLPS-ITER scenario. Using these inputs, in section 4.2 we show the effect of several different parameters that existed in DIV1D at the beginning of this project. The goal here is to show what one can and cannot achieve by tweaking those parameters. Finally, having established that the physics features which were then present in DIV1D could not account for the residual difference after tweaking the input parameters in DIV1D to the best of our ability, we describe the implementation of radial losses in DIV1D and the way that they influenced simulations results in section 4.3.

4.1. Paraphrasing SOLPS-ITER to DIV1D parameters

As a first step to comparing SOLPS-ITER with DIV1D a SOLPS-ITER run was identified that looked qualitatively compatible with DIV1D. Qualitatively compatible, here, was taken to mean: having similarly

shaped parallel profiles for the electron density and electron temperature. Such a run was identified as shot[7] #150662 (Wensing, 2020). This shot is one out of a series of simulations in which the gas puff was varied over the course of several shots and using a baffled[8] divertor. The geometry is TCV's. The simulation was run using SOLPS-ITER version 3.000.006 with drifts due to the electric and magnetic fields enabled (Schneider *et al*, 2006). Deuterium ions constitute the main plasma species, while the charged states of carbon constitute the impurities.

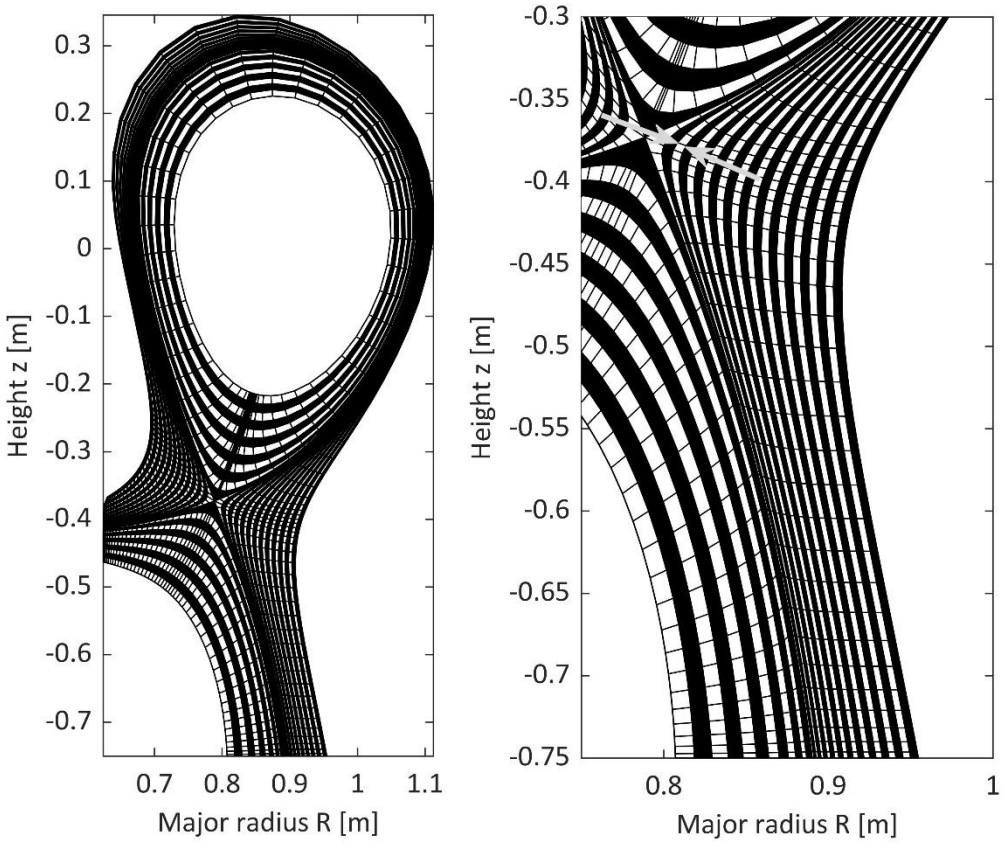

Figure 4.1. SOLPS-ITER (B2.5) grid of TCV. The grid is field-aligned and adjacent flux tubes have been coloured black or white alternatingly. The X-point, located near $(0.8, -0.37)$, is clearly visible. The separatrix, i.e. the position where the flux tubes go from closed to open just outside the last closed flux surface (LCFS), is located on the interface of the black and white flux tube that form the X-point. The scrape-off layer is the portion of the grid that starts out from the inner divertor (on the very left of the left-hand figure), then rounds the main chamber and terminates on the outer divertor (bottom). The two more crudely sampled portions of the grid are in the main chamber (which does not end up at the divertor) region and the private flux area (which does end up at the divertor). A zoomed-in version of the outer divertor is shown on the righthand side. The flux tube selected in the present section is wedged between two grey arrows (near the X-point).

As a starting point to the comparison, a flux tube needed to be identified that could serve as the basis of comparison. The flux tubes of SOLPS-ITER are indicated alternatingly in black and white in figure 4.1. A flux tube was selected that sat just outside the separatrix owing to the fact that the parallel heat flux was almost at its highest there. A flux tube slightly outside the separatrix was chosen because the heat flux changed rather rapidly as a function of the radial coordinate right at the separatrix. With the knowledge of today perhaps this was not the wisest choice, because outwardly pointing radial transport

[7] Although the word 'shot' is typically reserved for an experimental discharge, no such connotation should be attributed here. A 'shot' or 'shot number' relates to a database entry in an MDSPlus database that contains all the in- and output files of a SOLPS-ITER simulation. See MDSPlus (2020).

[8] A baffled divertor is a divertor configuration in which walls are placed between the divertor plasma and the main chamber to prevent leakage of neutrals and thus increase the neutral pressure in the divertor chamber.

(the direction of which is almost coincident with the direction of increasing major radius in figure 4.1, at least for the outboard divertor) seems to shift the strike point (the position of the peak heat flux at the divertor) to a position further removed from the separatrix, something which is identified in figure 4.2. The strike point constitutes the weakest link in the divertor because circumstances are most extreme there. If critical limits of materials engineering are exceeded anywhere, they are most likely exceeded there.

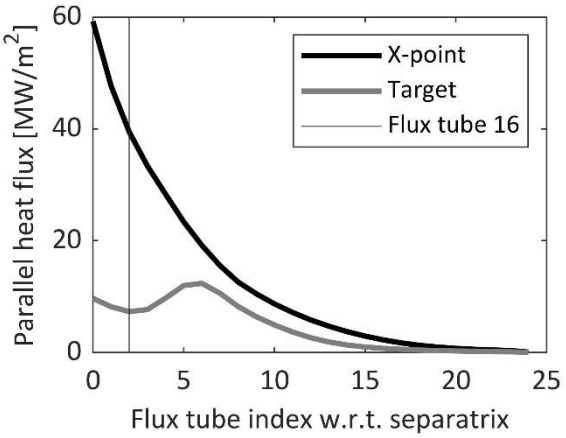

Figure 4.2. Footprint of the heat flux as a function of flux tube index. On the upstream side the parallel heat flux peaks closest to the separatrix while it peaks more deeply in the SOL at the target position. Note that this heat flux is parallel to the magnetic field, which sits at an angle of $\sin(\theta) = 0.065$ to the divertor target. It is bigger than the heat flux perpendicular to the divertor target by more than an order of magnitude. The heat flux perpendicular to the target, even at the peak position in the grey line, never exceeds $1\ \mathrm{MW/m^2}$. The flux tube selected in the present section is indicated by the vertical black line.

The separatrix in shot $\#150662$ is located at $i_y = 14$, where i_y is the index to the radial coordinate y. It can be readily observed in figure 4.1. Looking at the X-point one we identify two black and two white flux tubes that together form the X at the X-point. Focussing on the outer divertor leg (right side), the separatrix is the interface between the white flux tube to the right of the X-point and the black flux tube at the bottom of the X-point. The flux tube that was selected for this study is two positions to the right of that separatrix at $i_y = 16$. Determining the poloidal coordinate range is rather straightforward. There is a parameter in SOLPS-ITER that gives the position of the target, which is $i_x = i_{x,\mathrm{t}} = 96$ in this case. That gives the index to poloidal coordinate of the target side. The upstream position in DIV1D is taken to coincide with the X-point. In all fairness, this is imprecise language because the X-point only exists on the separatrix. Rather, a position had to be identified near the X-point on the flux tube in question, which is equally imprecise language for saying that a flux tube was picked by sight: this position is labelled $i_x = i_{x,\mathrm{u}} = 74$. Since all parameters of the 2PM are less than (inversely) linearly proportional to the connection length, one does not need worry strongly about being one cell too far up- or downstream.

Outside of these i_y and i_x which are (proxies for) the radial and poloidal coordinate, one more coordinate had to be determined to calculate this connection length, which is the pitch angle. After all, DIV1D is a code parallel to the magnetic field, which means that its main component is in the toroidal direction which is perpendicular to the plane of SOLPS-ITER. In terms of the necessary physics, this requires the simulator to obtain the pitch angle of the magnetic field. This, in turn, can be obtained by extracting the total magnetic field and the component thereof in the poloidal direction. Alternatively, a parameter exists in SOLPS-ITER that directly logs the distance along the magnetic field line. Either way, by requesting these or this parameter, it was possible to calculate the connection length to be $8.4\ \mathrm{m}$.

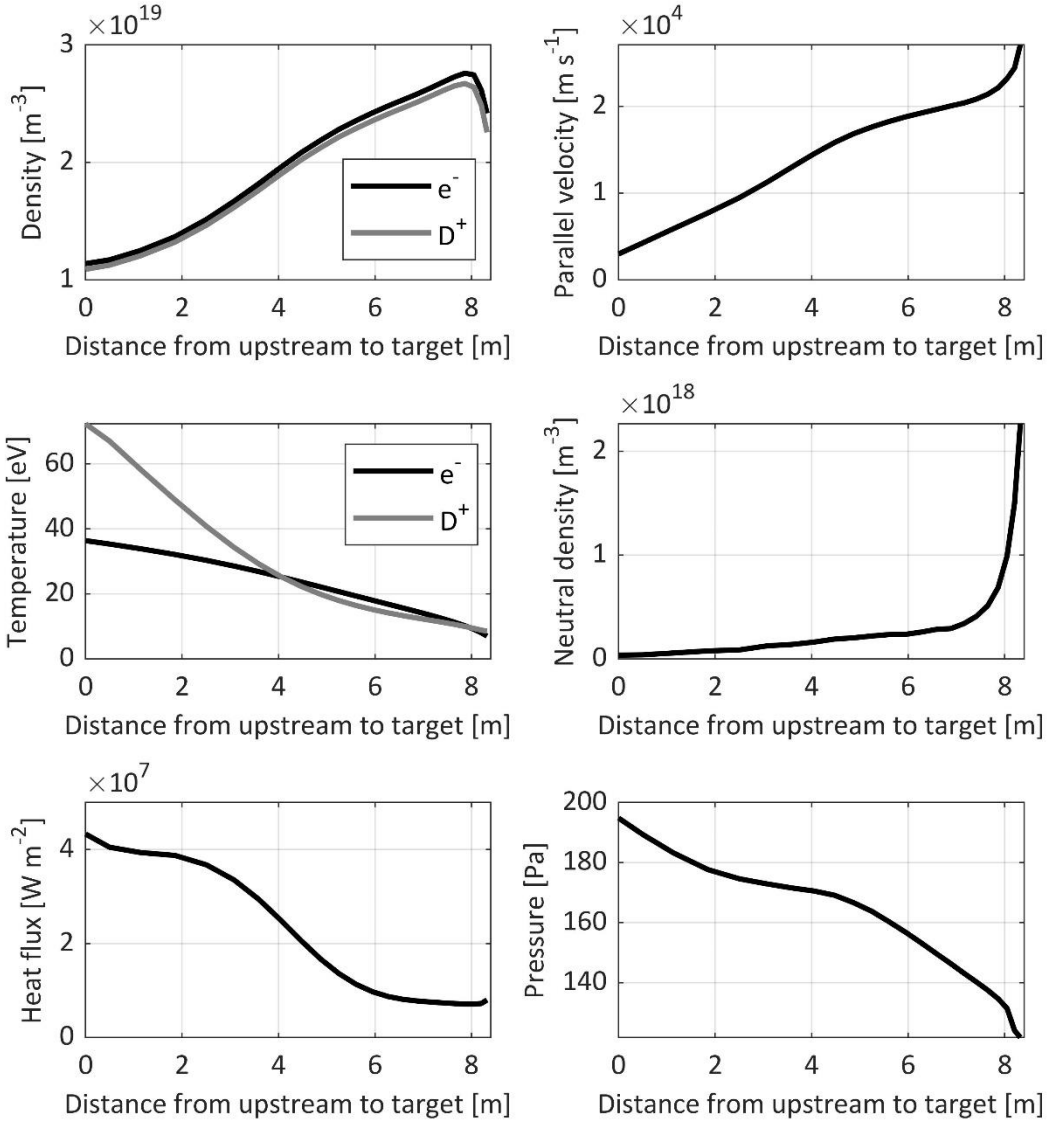

Figure 4.3. Several key plasma parameters from SOLPS-ITER shot #150662. While DIV1D assumes that $n_e = n_i$ and $T_e = T_i$, SOLPS-ITER keeps track of both. The density profiles are very close together but the temperature profiles do differ from ions to electrons. What draws the eye is the strong decline in the heat flux along the flux tube. These losses are incurred at rather high temperatures at which the DIV1D sources and sinks do not yet strongly come into play.

Having determined the start and end points as well as the length between them, it was then possible to start plotting parallel profiles of several key parameters. They have been plotted in figure 4.3. The electron and deuterium-ion densities were trivially obtained without prior knowledge of SOLPS-ITER, as were the electron and deuterium-ion temperatures, the parallel velocity, the heat flux and the pressure. The neutral deuterium density required a bit of care, however. As touched upon in section 2.2, SOLPS-ITER is actually a combination of two codes under the hood: B2.5 and EIRENE. The former is used for the plasma species while the latter is used for the neutrals. B2.5, however, also has a neutral model—one that basically treats neutrals as uncharged plasma species—which is numerically suppressed while B2.5 is coupled to EIRENE. The neutral density, for numerical reasons, is set to an exceedingly low but non-zero value in calculations. That does not mean, however, that the correct neutral density is automatically updated in the output files. One needs to 'manually' remember to get the EIRENE neutrals instead of the B2.5 neutrals. Trivial if you are a SOLPS-ITER expert, but not so if you are not.

As an intermezzo, it should also be noted that figure 4.3 shows that the simplification in DIV1D that $n_{\mathrm{e}} = n_{\mathrm{D}^+}$ seems valid enough. The simplification that $T_{\mathrm{e}} = T_{\mathrm{D}^+}$ seems to break down on the upstream side of the divertor leg, however. The question: Is that a problem? Probably not. The main effect of the temperature is via the heat flux. Looking at equation 2.15, there are two terms that make up the heat flux. Starting with the latter, we find that the conduction term is dominated by the electron heat conductivity which is more than a factor of 30 bigger than that of the ions. The convection term is not that important, either, because the divertor is typically conduction-limited which means that the convection term is much smaller than the conduction term. Underestimating the ion temperature will lead to an underestimation of the pressure in the divertor leg with respect to that found in SOLPS-ITER, but little else besides that.

Looking at figure 4.3 through the eyes of the 2PM, the six governing parameters are easily identified. One of them trivially so, because the connection length L was actually required to make these graphs and already discussed before. For completeness' sake, $L = 8.4\ \mathrm{m}$. The other two independent parameters are the upstream (X-point) density and parallel heat flux, which are $n_{\mathrm{u}} = 1.1 \times 10^{19}\ \mathrm{m}^{-3}$ and $q_{\parallel,\mathrm{X}} = 45\ \mathrm{MW\ m}^{-2}$. The three dependent parameters are the upstream and target temperature, and the target density. Those are $T_{\mathrm{u}} = 36\ \mathrm{eV}$, $T_{\mathrm{t}} = 7\ \mathrm{eV}$ and $n_{\mathrm{t}} = 2.4 \times 10^{-19}\ \mathrm{m}^{-3}$. The target temperature of $7\ \mathrm{eV} < 10\ \mathrm{eV}$ which means that this plasma is in the detached state according to Stangeby (2018). The two criteria set by Dudson *et al* (2019) with respect to achieving detachment are met likewise, as exemplified by the dual drop in heat flux and pressure between upstream and target visible in the bottom two graphs of figure 4.3. What stands out is the fact that the heat flux drops steeply halfway through the divertor leg—in a temperature regime that is too high for the power sinks in DIV1D to play any significant role yet.

Outside of the three independent 2PM parameters L, n_{u} and $q_{\parallel,\mathrm{X}}$ DIV1D needs several other input parameters to run. One of these is the pitch angle θ that enters into the neutral diffusion equation 2.19 as $D \propto \sin^2\theta$. It is visualised in figure 2.1. As neutrals in DIV1D only really play a significant role in the direct vicinity of the target, the pitch angle for DIV1D was taken to coincide with the pitch angle in SOLPS-ITER at the target. Its value is equal to 0.06. Recycling was left at 100% in DIV1D, which reflects that the high-recycling regime (which is synonymous to the conduction-limited regime that is assumed in working through the 2PM) is a necessary prerequisite for the transition into detachment (Stangeby, 2000). Finally, there is the carbon concentration, which is taken into account entirely differently in SOLPS-ITER. Using B2.5, SOLPS-ITER solves the complete balance equations for all six charge states of carbon. DIV1D, meanwhile, simply assumes a constant impurity fraction and goes from there. To define an impurity fraction for use in DIV1D the charge-weighted sum of all charge states' populations was calculated according to

$$\xi'_Z(x) = \frac{1}{n_{\mathrm{e}}(x)} \sum_{Z=1}^{6} Z \times n_{\mathrm{C}^{Z+}}(x). \tag{4.1}$$

where Z is the charge number, $n_{\mathrm{C}^{Z+}}$ is the position-dependent density of carbon with charge number Z, $n_{\mathrm{e}}(x)$ is the electron density and $\xi'_Z(x)$ is the impurity fraction. To go from 4.1 to a flux-tube constant value as is currently employed in DIV1D, one additionally needs to integrate 4.1 over the divertor leg

$$\xi_Z = \frac{1}{L} \int_0^L \xi'_Z(x)\mathrm{d}x. \tag{4.2}$$

Equations 4.1 and 4.2 have been plotted together in figure 4.4, where 4.1 has been labelled 'SOLPS-ITER' while 4.2 has been labelled 'DIV1D'. What stands out from the position-dependent variant is the sharp peak that exists in the direct vicinity of the divertor target. Where values through the divertor are limited to about 0.01 (or 1%) they shoot up to almost 0.05 (or 5%) directly in front of the target. This does not have much effect on the divertor-leg average as becomes apparent from how close the black and grey lines lie together over much of the flux tube. The effect of this peak directly in front of the target has not been taken into consideration in this thesis but it could be added to the DIV1D code without too much trouble in principle. It is expected that this peak might have a very limited effect in the range where the plasma temperature has dropped below $10\ \mathrm{eV}$ half a meter in front of the target—which is the temperature range in which carbon starts to radiate significantly (Post *et al*, 1977) and as a result DIV1D might be seen to transition into a fully detached state marginally more quickly. The effect is expected to be minor though. And it definitely cannot account for the strong loss of power in the middle of the divertor leg however. The carbon concentration there is actually a bit below that in DIV1D.

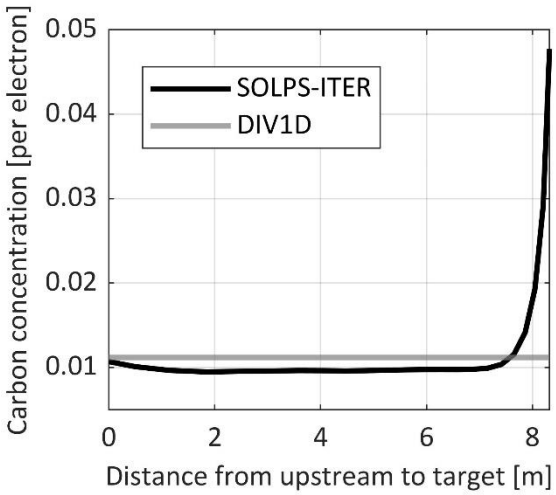

Figure 4.4. Carbon concentration as a function of the parallel distance from upstream to target. SOLPS-ITER (black) has a carbon concentration that is rather constant across much of the flux tube. It rises slightly towards the upstream and it rises very strongly near the target. While it stays in the ballpark of 1% on the upstream, it rises to as much as 5% near the target. Taking the average over the length of the flux tube, the carbon concentration in each cell weighted according to that cell's parallel length, an input parameter for DIV1D (grey) can be obtained. It sits close to the constant value found throughout much of the SOLPS-ITER flux tube which reflects that the peak near the target, although large, is spatially limited and thus has limited effect on the average.

4.2. DIV1D without radial losses

Having extracted the relevant parameters from the SOLPS-ITER simulation, we could then try to reproduce these results with DIV1D. The initial approach was to just input the obtained parameters from section 4.1 and start simulating. However, serving those parameters to DIV1D as inputs caused the code to fail on multiple occasions. After repeated attempts with different parameters failed to yield converging results, a different strategy was employed which made good use of a feature in DIV1D that allows the user to serve it with a file called `div1d_restart_old.txt` at the beginning of a simulation. This file contains the information of a previously converged result and can be used to initialise all parameters on the grid. It is meant to achieve more speedy convergence and is particularly helpful with parameter scans. In such scans, typically, one parameter is incrementally changed over a series of runs. Each simulation, along with a results file that contains many different plasma parameters sampled at a fixed frequency, also outputs the file `div1d_restart_new.txt` which contains full profiles of all necessary

parameters at the end of the last timestep. One could expect that the plasma scenario looks rather similar between two runs of a parameter scan, and that makes it useful to serve the previous run's output as the initial condition to the next run.

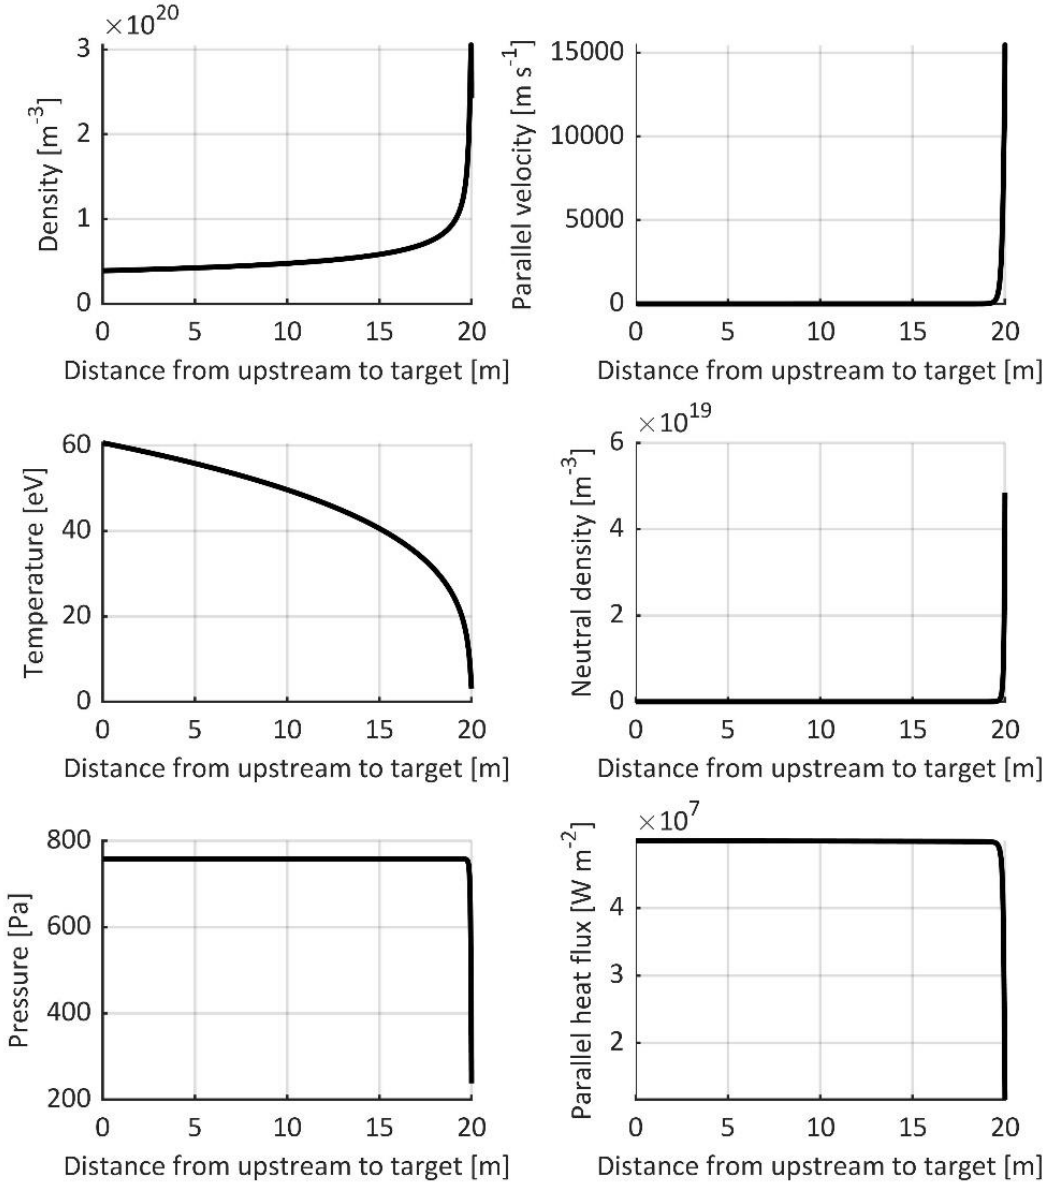

Figure 4.5. Simulation of a deuterium plasma in DIV1D as part of a density scan by Westerhof (unpublished). The shot name is 3.9. Graphs are of the plasma density, the parallel velocity, the temperature, the neutral density, the pressure and the parallel heat flux as functions of the distance from upstream to target. The upstream density is $3.9 \times 10^{19}\ \mathrm{m}^{-3}$, the upstream heat flux is $50\ \mathrm{MW/m^2}$ and the connection length is $20\ \mathrm{m}$. Additionally, wall recycling is at 100%, the carbon concentration at 1% and the pitch angle $\sin\theta = 0.1$. The neutral residence time is numerically infinite (10^{20} s). The plasma is detached: $T_\mathrm{t} = 3.06\ \mathrm{eV} < 10\ \mathrm{eV}$.

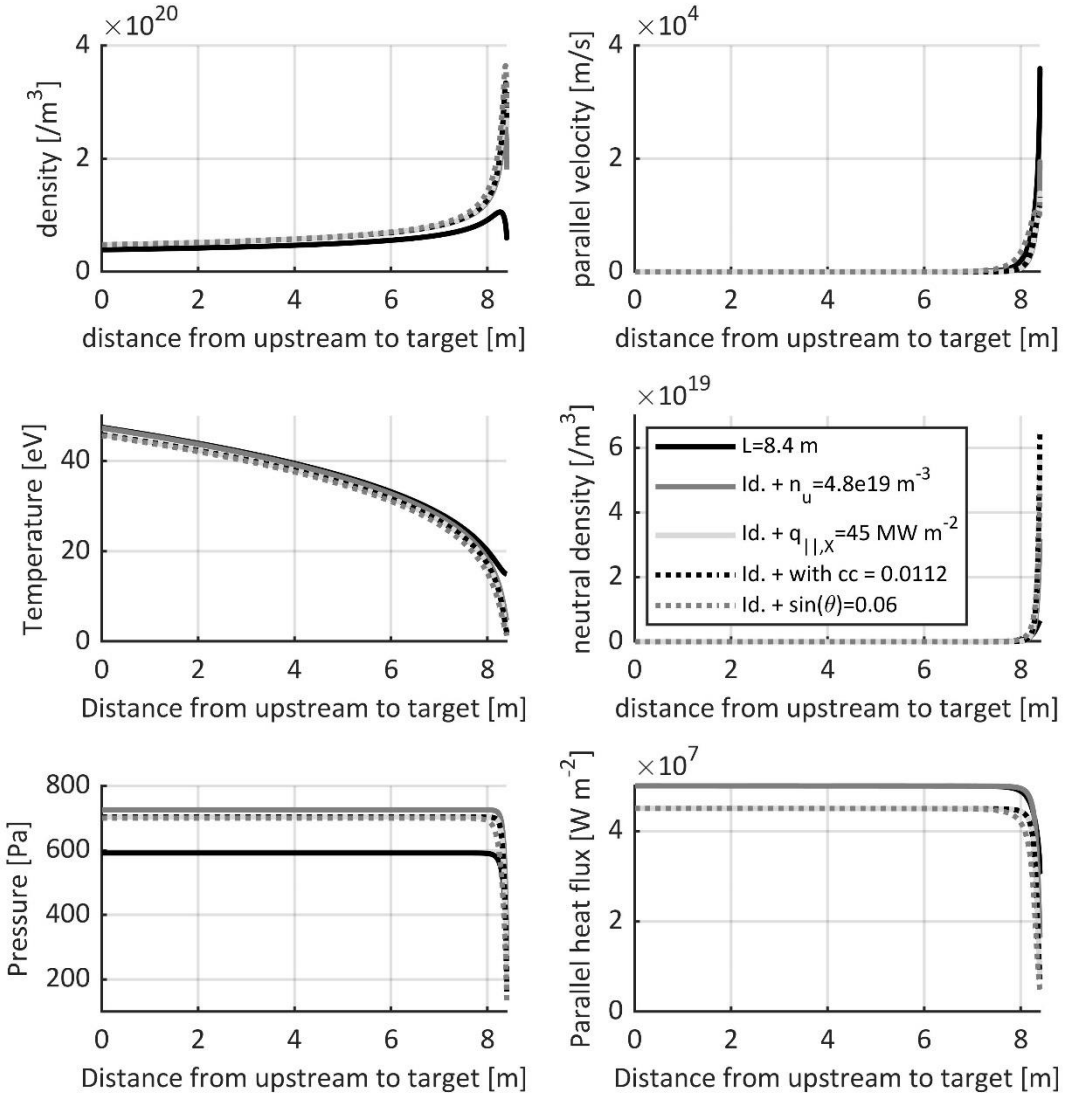

Figure 4.6. Starting out from the DIV1D shot 3.9 parameters were adapted to reflect SOLPS-ITER shot #150662. First the connection length was retracted from $20\ \mathrm{m}$ to $8.4\ \mathrm{m}$. Initially this was done while keeping the upstream density unchanged at $3.9 \times 10^{19}\ \mathrm{m}^{-3}$. This scenario is plotted as the solid black line. However, as this significantly increases the target temperature, the upstream density was then increased to an upstream density of $4.8 \times 10^{19}\ \mathrm{m}^{-3}$; it is plotted as the solid dark grey line. Following this, the upstream heat flux $q_{\parallel,\mathrm{X}}$ was reduced from its initial value of $50\ \mathrm{MW\ m}^{-2}$ to $45\ \mathrm{MW\ m}^{-2}$. The effect of this has been plotted in solid light grey. It overlaps almost completely with the solution obtained after changing the carbon concentration from 1% to 1.12%, indicated by the dotted black line. A reduction in upstream heat flux leads to a slight reduction in the upstream temperature in keeping with the two-point model and a relatively strong decrease in the target temperature. As a result, the temperature shifts to slightly lower values across the entire profile. This logically translates into a slight drop in the pressure. The density at the target, meanwhile, increases. Finally, the change in the pitch angle leads to the graph plotted as the dotted dark grey line. It has the effect of decreasing the peak neutral density slightly, which is logical since a smaller pitch angle increases the speed of diffusion of neutrals in the flux tube.

Using this option to our advantage, we started with a converged result from a previously run density scan. The $3.9 \times 10^{19}\ \mathrm{m}^{-3}$ scenario—which was also the starting point for the grid convergence and artificial viscosity scan—is plotted in figure 4.5. Although many parameters are different in that simulation—the connection length is a factor three longer, and the upstream density more than a factor three higher—the target temperature is about where it should be. It is $3.1\ \mathrm{eV}$, which brings it in the

range of plasmas that have entered the detached regime below $10\ \mathrm{eV}$ but which still has a detachment front that is localised very near to the target, similar to the SOLPS-ITER $\#150662$ shot. Then, one by one the key parameters were adapted to reflect those obtained in section 4.1. The cumulative effect of several changes is displayed in figure 4.6.

First, the correct connection length was set to its SOLPS-ITER value of $8.4\ \mathrm{m}$. At an upstream density of $3.9 \times 10^{19}\ \mathrm{m}^{-3}$ this had the result of strongly increasing the target temperature, which had been in the correct ballpark previously. Decreasing the divertor leg has a positive effect on the upstream temperature; it is still higher than the SOLPS-ITER value but it decreased by 25% with respect to its initial value in keeping with the 2PM. To drop the target temperature back to within the $0.1\ \mathrm{eV} < T_\mathrm{t} < 10\ \mathrm{eV}$ range the upstream density was increased to $4.8 \times 10^{19}\ \mathrm{m}^{-3}$. The resulting parallel profiles are given by the solid dark grey line in figure 4.6. However, this came at the cost of increasing the upstream density even further; and that while it was already more than a factor three too high. Physics wise, the density acts on almost all the sinks in DIV1D so requiring a higher density actually reflects that DIV1D sinks are not effective enough at the densities obtained in SOLPS-ITER.

Next, the parallel heat flux was reduced from $50\ \mathrm{MW\ m^{-2}}$ to $45\ \mathrm{MW\ m^{-2}}$ at the upstream. This led to a lower temperature along the entire divertor leg as well as an increase in the target plasma density. These findings are also in keeping with the 2PM. The resulting parallel profiles are given by the light grey line in figure 4.6. The ensuing change in carbon concentration from 1% to 1.12% had little effect on the solution. A change, finally, to reflect the correct pitch angle θ has been plotted as the grey dotted line. This parameter enters into the diffusion coefficient for the neutrals as $D \propto \sin^{-2}\theta$. This was visually illustrated in figure 2.1: increasingly smaller value for θ put the direction of plasma transport at a increasingly larger angle with respect to the direction of neutral transport. As a result, neutral transport picks up in the flux-tube frame-of-reference. This results in a slightly broader neutral peak directly in front of the target and, as a result, a slightly broader plasma density peak owing to the broadening of the ionisation layer. Consequently, the ionisation layer and resultant power flux drop also broaden slightly, which finally causes the sharp decline in the temperature that typically happens near the target to move marginally further upstream. As figure 4.6 shows, however, the effects are very small.

Outside of the parameters named up to now, there are three additional parameters that can be changed. One could still think to change the recycling coefficient, or one could turn the knobs of either the neutral residence time or the gas puff. They all involve the neutrals, just as the pitch angle did. As with the pitch angle, however, playing with these parameters turned out to only have limited effect. Refer to figure 4.7. Decreasing the level of recycling caused the density peak directly in front of the target to become lower. This is logical considering that this peak is linked to the ionisation peak, which depends on the number of available neutrals. That number of available neutrals, in turn, is dependent on the number of neutrals that is reinserted locally via wall recycling. Decreasing the level of recycling also created a non-zero positive parallel velocity at the X-point, reflecting that particles that do not return into the flux tube via the wall need to be inserted on the upstream side of the domain to maintain the particle balance. The effect of changing the neutral residence time—a parameter that dictates over what typical time scale τ_neutral the neutral inventory of the flux tube is lost, see equation 2.20—has a similar effect. Although it is in principle a flux-tube–wide loss the effect in DIV1D is much the same as in the case of recycling because, as figure 4.5 shows, there is no appreciable neutral density except right near the target. It should come as no surprise, then, that the effect of a non-infinite neutral residence time is to lower the density peak near the target and to create a positive in-flux of plasma particles at the X-point, which again requires a non-zero positive parallel velocity at the X-point. Tweaking the amount of gas puff, finally, has the oppositive effect: it creates a surplus of neutrals along the divertor

leg, which results in a slightly higher plasma density across the board and a higher neutral content directly in front of the target, while it causes the parallel velocity to drop below zero near the X-point to siphon away excess particles.

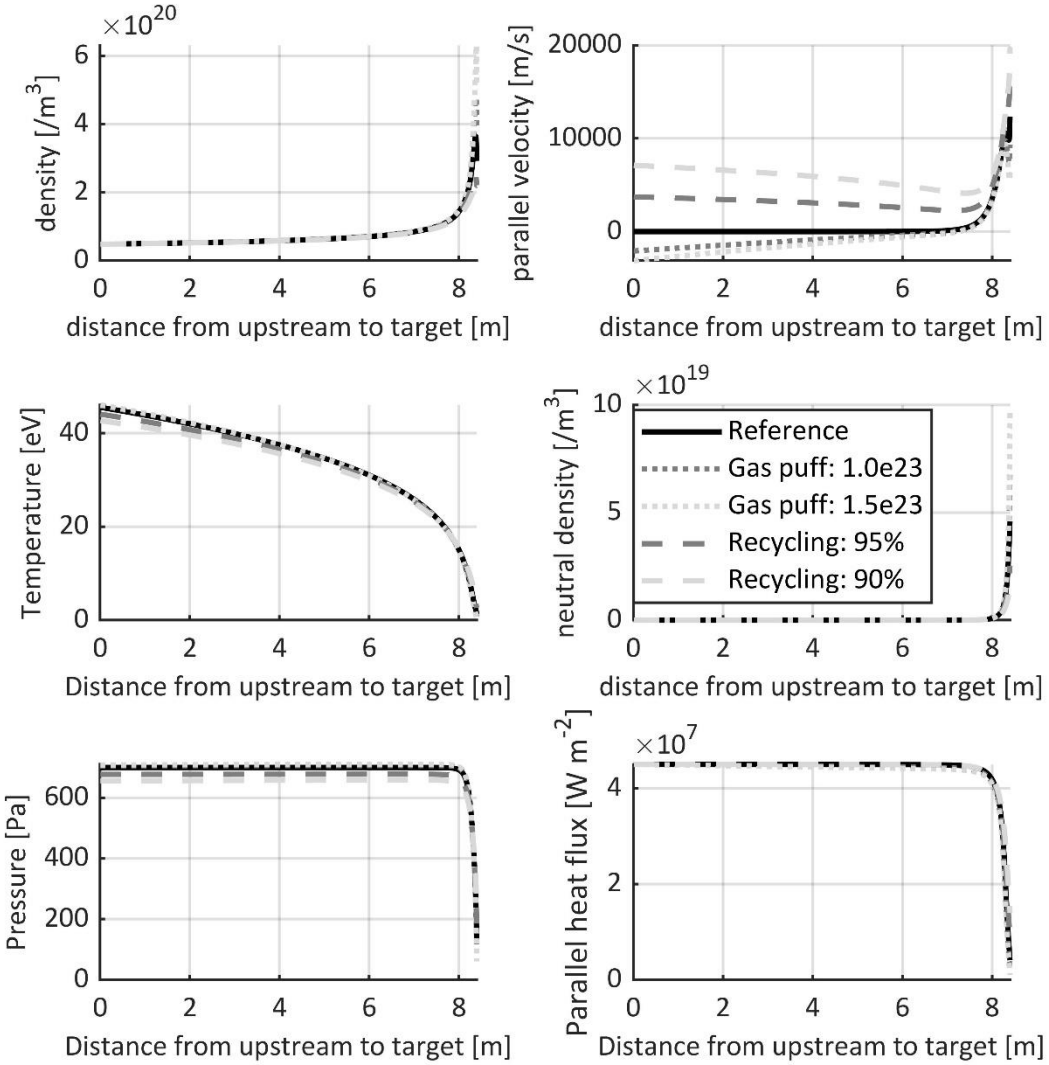

Figure 4.7. Solutions of DIV1D simulations in which either gas puff or recycling were added to a reference simulation. This reference simulation is the last simulation displayed in figure 4.6 (titled "Idem with $\sin(\theta) = 0.06$"). Gas puff and recycling work similarly but in opposite directions. The former increases the target neutral and plasma densities while causing the upstream parallel velocity to decrease (below zero), the latter *vice versa*.

Having tweaked all possible parameters in DIV1D, it seemed that there was no way to obtain the result of SOLPS-ITER in the code such as it existed at the time. Ultimately, this made sense considering the 2PM. After all, according to that model the key parameters that dictate the physics in the SOL are the parallel heat flux, the upstream density and the connection length. Those in turn determine the upstream and target temperatures as well as the target density. Since all these parameters were fixed going into the comparison, there was little to tweak. And indeed, while other parameters—such as the recycling coefficient, the pitch angle and gas puff—are available they only proved to have limited effect on the solutions. This understanding paved the way for something new to be included in the DIV1D code: radial losses.

4.3. DIV1D with radial losses

An inventory of the different loss channels suggested that a primary heat loss channel for SOLPS-ITER was not included in DIV1D. As these losses take place perpendicularly to the flux tube that is modelled in DIV1D, there is no way to self-consistently model the underlying perpendicular heat fluxes. To work around this limitation, an attempt was made to define effective perpendicular loss terms to the DIV1D code. These changes were implemented through a new subroutine calculate_radial_losses that defines one of two possible effective perpendicular sink terms. The source code has been included as Appendix A. The most straightforward implementation is to spread the losses evenly along the flux tube by defining an effective perpendicular sink

$$S_\perp = -\eta_\perp \frac{q_{\parallel,\mathrm{x}}}{L}, \tag{4.3}$$

where $\eta_\perp$ is the fraction of the parallel heat flux $q_\parallel$, in units of $\mathrm{W\ m^{-2}}$, that is dissipated perpendicularly to the flux tube. By dividing through the divertor leg length L this can be rewritten into an effective volumetric sink $S_\perp$, in units of $\mathrm{W\ m^{-3}}$.

Motivated by the fact that the perpendicular losses in SOLPS-ITER were primarily localised about halfway along the flux tube from X-point to target, a potentially more truthful formulation could include a position dependence. To this end, a second way to define an effective loss term was formulated as

$$S_\perp(x) = -\frac{\eta_\perp q_{\parallel,\mathrm{x}}}{\sqrt{2\pi(\Delta r)^2}} \exp -\frac{1}{2}\left(\frac{x - x_0}{\Delta r}\right)^2, \tag{4.4}$$

which is very similar to 4.2 except that the factor of L^{-1} has been replaced by a Gaussian centred around x_0 with a 1-over-e width of Δr. Note that the total amounts of particles lost perpendicularly in equations 4.3 and 4.4 are equal to within 5% percent for $|L - x_0|/\Delta r \leq 2$ and equal to within less than 1% for $|L - x_0|/\Delta r \leq 3$. This is the formulation that was used to run the simulations in this section. It is not, however, the version that currently exists in DIV1D. It was later decided that it made more sense to calculate the prefactor by numerical integration of $S_\perp(x)$ over L so that the total loss of power would be equal to $-\eta_\perp \times q_{\parallel,\mathrm{x}}$ always.

Scans with variable radial loss fraction $\eta_\perp$ were run using equations 4.3 and 4.4 respectively. Through a concurrent linear increase of the radial loss fraction and a linear decrease of the upstream density at the X-point, a series of stable solutions were obtained for both formulations of the radial heat loss term. Differences between the parallel-position–dependent formulation and the constant formulation turned out to be slim. This makes sense when taking into the sources that play a role in DIV1D's heat balance. By inspection of the graphs, one can see that there is only a significant neutral population in the direct vicinity of the target. This precludes any significant contribution of charge exchange or excitation losses, which go linear in the neutral density. The effect of the other two loss channels is similarly limited. Recombination effects are negligible above $1\ \mathrm{eV}$, and impurity radiation peaks at $7\ \mathrm{eV}$. Inspecting the temperature graph, one can see that temperatures are higher than that everywhere except, again, in the direct vicinity of the target. This section describes these simulations in more detail.

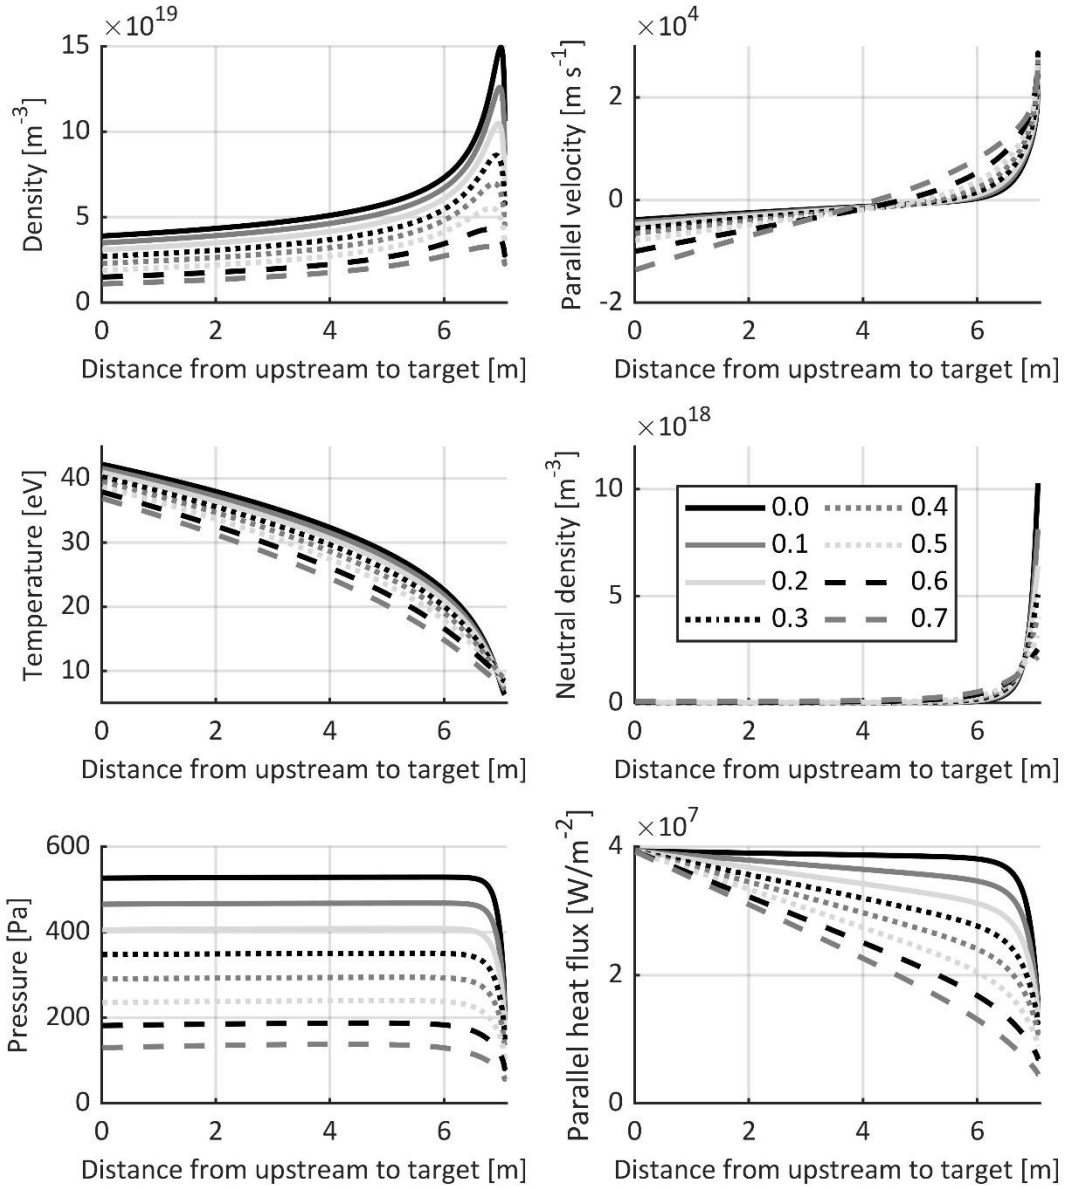

Figure 4.8. Parameter scan in which the radial loss factor was varied in tandem with the upstream density in the case of linear losses. Increasing the loss over the flux tube from $\eta_{\mathrm{r}} = 0.0$ (0%) to $\eta_{\mathrm{r}} = 0.7$ (70%) of the incoming parallel heat flux at the X-point (located at 0 m) allows for the upstream density to be dropped at constantly spaced increments of $\Delta n_{\mathrm{u}} = 0.4 \times 10^{19}\ \mathrm{m}^{-3}$ starting from $n_{\mathrm{u}} = 3.9 \times 10^{19}\ \mathrm{m}^{-3}$ at $\eta_{\mathrm{r}} = 0.0$ to $n_{\mathrm{u}} = 1.1 \times 10^{19}\ \mathrm{m}^{-3}$ at $\eta_{\mathrm{r}} = 0.7$. All simulations were performed for a flux tube of $7.065\ \mathrm{m}$ with an incoming heat flux at the X-point of $39.35\ \mathrm{MW/m^2}$. The carbon concentration is 1%, recycling set to 100% and the gas puff source is set to $1.5 \times 10^{23}\ \mathrm{s}^{-1}$. This gas puff is equally distributed over the divertor leg, which amounts to a particle source of $0.21 \times 10^{23}\ \mathrm{m}^{-3}\ \mathrm{s}^{-1}$.

As discussed, DIV1D comes packed with the possibility to feed an output file from an earlier run as an input file to the next one. By starting out from a converged run and increasing or decreasing any parameter of interest by tiny steps over the course of several runs, one can slowly transition from one regime into another. In the case at hand, it was possible to move from a converged run with an upstream density of $n_{\mathrm{u}} = 3.9 \times 10^{19}\ \mathrm{m}^{-3}$ with $\eta_{\perp} = 0.0$ to an upstream density of $1.1 \times 10^{19}\ \mathrm{m}^{-3}$ with $\eta_{\perp} = 0.7$ while keeping the target temperature in the range between 0.1–$10\ \mathrm{eV}$. Simulations were run for a parallel heat flux at the X-point of $q_{\parallel} = 39.35\ \mathrm{MW\ m}^{-2}$ and a connection length of $L = 7.065\ \mathrm{m}$. The

carbon concentration was set to 1%. This set of parameters is slightly different from those presented in section 4.1 because the upstream parameters taken from SOLPS-ITER had been taken from a position a bit past the X-point in the direction of the target; this was corrected in section 4.1 but these simulations were not repeated alongside. Figure 4.8. shows the simulation results for these parameters using constant radial losses.

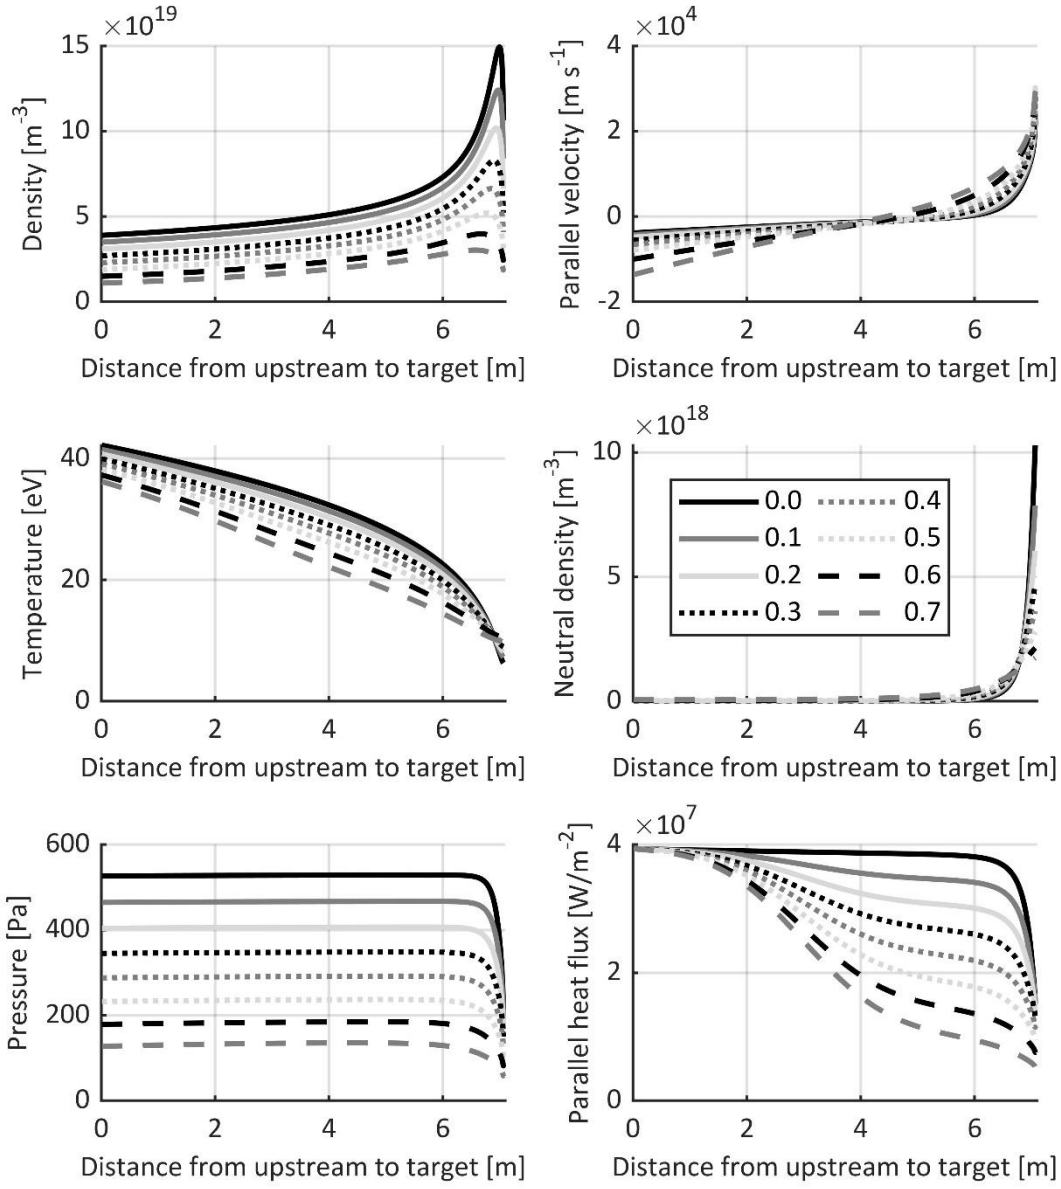

Figure 4.9. Parameter scan in which the radial loss factor was varied in tandem with the upstream density in the case of a gaussian radial loss profile centred at 3 m with a width of 1.167 m. All other simulation parameters are identical to those of the linear losses variant. This scan was performed using a legacy version of the radial losses in which the normalisation factor of the gaussian was still calculated analytically. This resulted radial losses that were slightly less extensive than the legend suggests as the tails of the distribution were cut off at the domain edges while the normalisation factor did include them. This error is less than 3%.

What can be seen in figure 4.8 is that increasingly stronger radial losses drive down the parallel heat flux along the flux tube. The temperature is rather indifferent to this change. This is logical from the

perspective of the (extended) two-point model. It was already established in section 2.4 that the upstream temperature is rather impervious to changes in any parameter. For the target temperature, matters are slightly more nuanced. The target temperature, according to the extended two-point model, follows $T_\mathrm{t} \propto \left(1 - f_\mathrm{pow}\right)^2$ which means that larger loss fractions quadratically relate to lower target temperatures. However, the basic two-point model teaches that $T_\mathrm{t} \propto n_\mathrm{u}^{-2}$. Since the radial power loss was varied proportionally to the upstream density, i.e. $n_\mathrm{u} \propto (1 - \eta_\perp)$, and total power loss fraction is dominated by the radial power loss, i.e. $f_\mathrm{pow} \approx \eta_\perp$, it follows that the target temperature $T_\mathrm{t} \propto n_\mathrm{u}^{-2}\left(1 - f_\mathrm{pow}\right)^2 \approx (1 - \eta_\perp)^{-2}(1 - \eta_\perp)^2 = 1$. The target temperature should thus not vary strongly over this scan either.

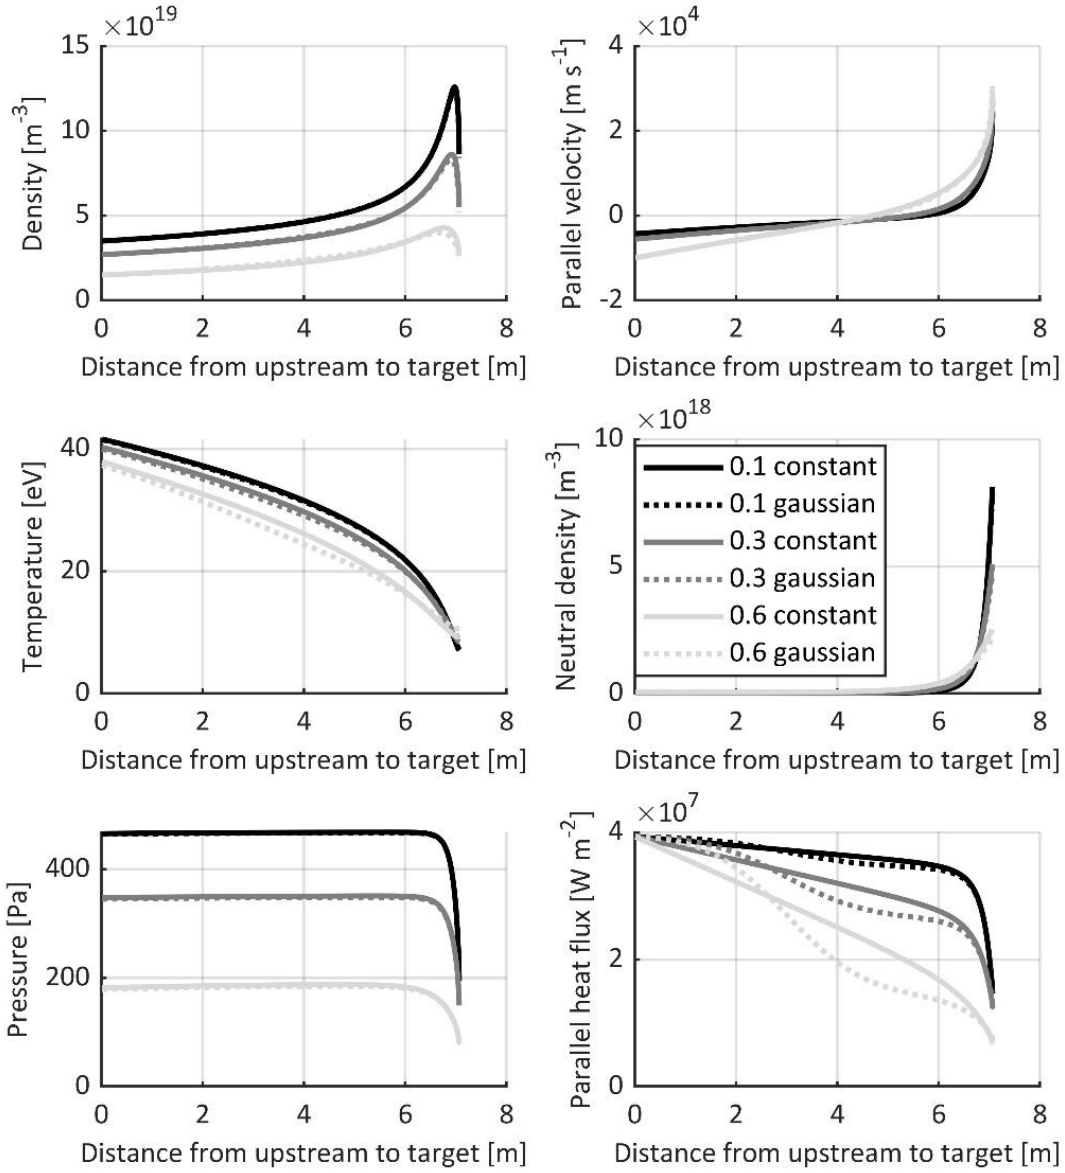

Figure 4.10. Comparison of constant versus gaussian radial losses for three different loss fractions $\eta_\perp = 0.1, 0.3, 0.6$. The bottom-right graph clearly shows differences between the heat flux profiles for any given $\eta_\perp$. However, the differences in these heat flux profiles have almost no effect on the other parameter profiles.

Redoing the scan for the same parameters but now using gaussian rather than constant losses yielded the simulation results displayed in figure 4.9. Looking at the bottom-right graph first we immediately identify the difference in the parallel heat flux. Where the heat was first lost at a constant rate over the flux tube there is now a clear peak in the losses. Very little heat is removed over the first meter of the flux tube as compared to the constant loss scenario. It then picks up quickly, the heat flux reaching its steepest decline at around $3\ \mathrm{m}$, which coincides with the position at which the gaussian was centred. Beyond that point the curve starts to flatten. This heat flux profile is very similar to that of $\#150662$. The effect on the other parameters is, interestingly, comparable to what was discussed for the constant-loss scenario. This is more clearly visible in figure 4.10, in which the constant loss scenario has been plotted against the gaussian loss scenario for several different values of the radial loss fraction $\eta_\perp$. In essence, this result is nothing but a confirmation of what the extended 2PM says. After all, in formulating a power loss fraction f_{pow} between the upstream and target sides it leaves open the question of where along the way these power losses are incurred. As long as it is lost somewhere between the X-point and the target parameters would be expected to behave the same. This is exactly what figure 4.10 shows.

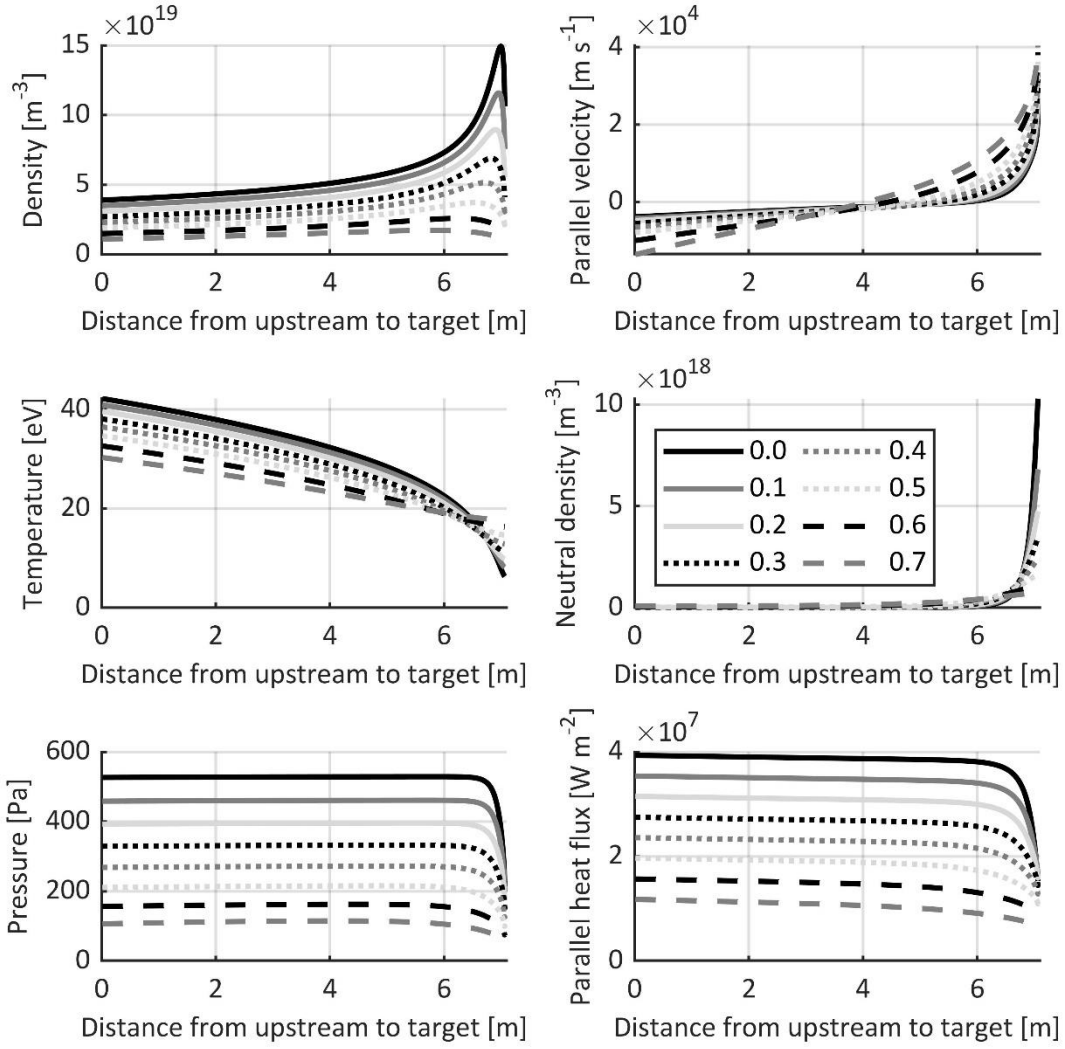

Figure 4.11. Parameter scan of an alternative approach to reducing the heat flux by taking into account the perpendicularly lost fraction directly at the X-point. Ignoring the perpendicular losses along the flux tube and instead making the upstream heat flux $\eta_\perp$-dependent, we set the heat flux to $q_{\parallel,\mathrm{X}}(\eta_\perp) = (1-\eta_\perp)q_{\parallel,\mathrm{X}}(0.0)$ where $q_{\parallel,\mathrm{X}}(0.0) = 39.35\ \mathrm{MW\ m^{-2}}$ similarly to before. The resulting solutions deviate significantly from those obtained via constant and gaussian radial losses. The upstream temperature decreases more strongly over the course of the parameter scan, as does the target density. The target temperature, meanwhile, increases.

One could suggest, at this point, that rather than siphoning off the heat flux at a given point in the flux tube, one might as well limit the heat flux that entered the flux tube in the first place. Rather than choosing a certain $q_{\parallel,\mathrm{X}}$ and $\eta_\perp$, one could consider doing away with the perpendicular losses entirely and simply setting the upstream heat flux to $(1-\eta_\perp)\times q_{\parallel,\mathrm{X}}$. This tactic was adopted in figure 4.11 There are notable differences compared to figures 4.8 and 4.9. The upstream temperature drops much faster than previously, and so does the target density. The target temperature, meanwhile, rises to much higher values than previously. So even though the exact loss profile is not very important, it is paramount that the heat does make it to the beginning of the flux tube. This can also be understood in the context of the (extended) 2PM: the exponents of $q_{\parallel,\mathrm{X}}$ and $(1-f_{\mathrm{pow}})$ are identical neither in the scaling for the upstream temperature, for the target temperature or for the target density.

One issue that persisted throughout this project was that the implementation of the perpendicular losses in DIV1D broke down in fully detached scenarios, i.e. scenarios in which the detachment front had broken loose from the target and receded towards the upstream. This was tentatively linked to an unphysical negative heat flux at some point in the flux tube, which would cause the heat to start flowing from the target towards the upstream. However, as there was no heat left there, this was impossible. The problem was that an attempted fix to this issue did not yield satisfactory results. Finally, at the eleventh hour we managed to debug the code—which had previously kept crashing once the plasma breached the threshold to full detachment, i.e. when the temperature hit $0.1\ \mathrm{eV}$ at the target—so that solutions were also obtained in a regime where the plasma was fully detached. Figure 4.12 shows an example of such a fully detached plasma ($L=7.065$, $n_{\mathrm{u}}=2.7\times10^{19}\ \mathrm{m^{-3}}$, $q_{\parallel,\mathrm{X}}=39.35\ \mathrm{MW\ m^{-2}}$ and constant losses with $\eta_\perp=0.5$). The way that the bug fix was implemented that enabled this can be found in box A.2 of appendix A.

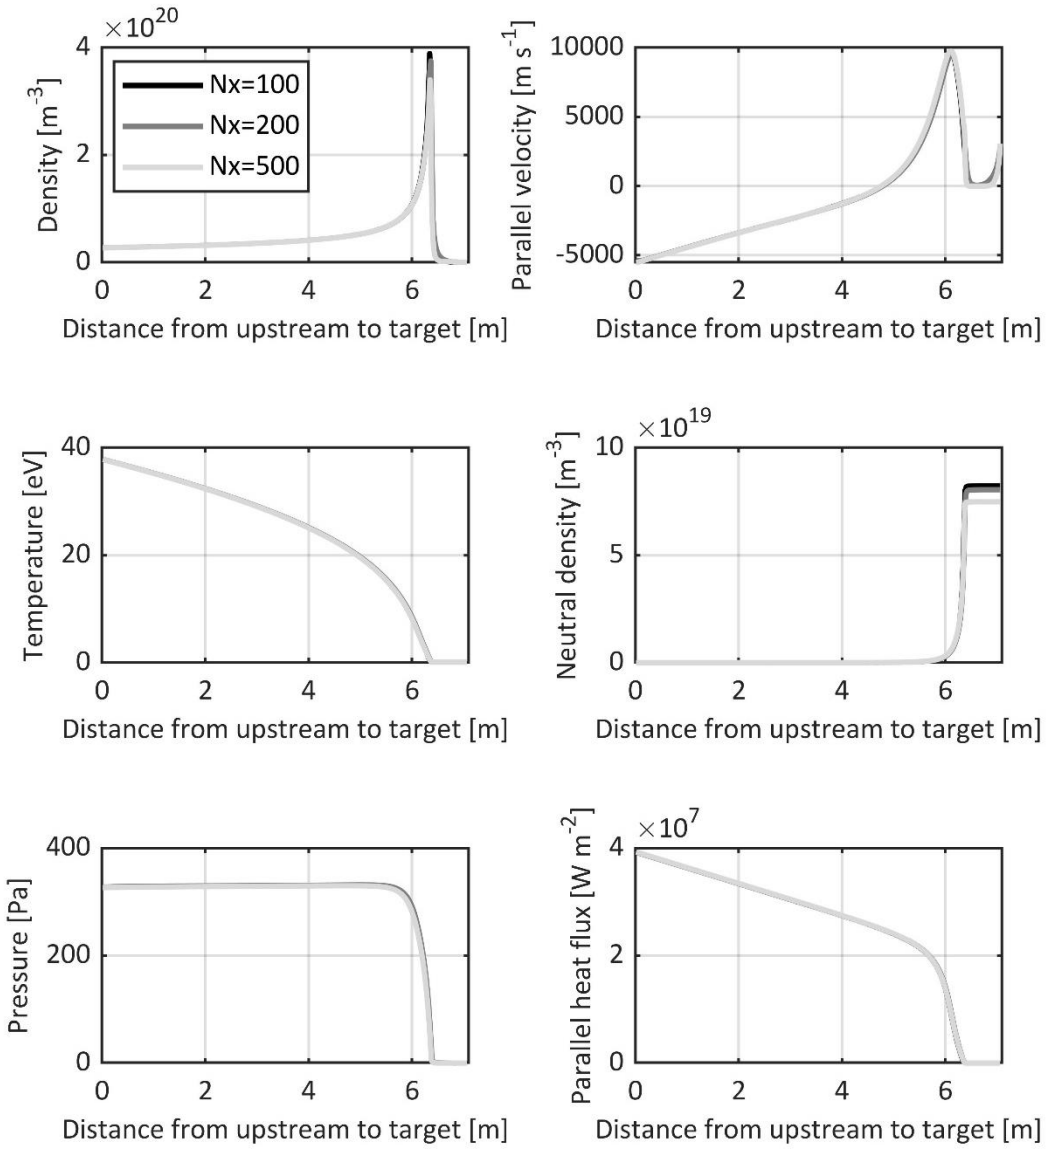

Figure 4.12. A fully detached plasma using parameters $L = 7.065$, $n_{\mathrm{u}} = 2.7 \times 10^{19}\ \mathrm{m}^{-3}$, $q_{\parallel,\mathrm{X}} = 39.35\ \mathrm{MW\ m}^{-2}$ and constant losses with $\eta_{\perp} = 0.5$. Simulations were performed for 100-cell, 200-cell and 500-cell grids. Gas puff is turned on at $1.5 \times 10^{23}\ \mathrm{s}^{-1}$. That the plasma is fully detached can be observed from the fact that the temperature has dropped to 0.1 eV at the target (which is the lowest possible value for reasons of numerical stability) while the ionisation peak has receded about half a meter from the target. Beyond the detachment front the neutral density quickly rises to a high and constant value up to the target. Without perpendicular losses, the detachment front would not have receded into the divertor. Stronger still, even at $\eta_{\perp} = 0.4$ the plasma could not access full detachment. This can be seen in figures 4.8–4.12 where the $\eta_{\perp} = 0.4$ scenario was plotted at the same density of $n_{\mathrm{u}} = 2.7 \times 10^{19}\ \mathrm{m}^{-3}$.

5. ELM-induced dynamics

This chapter discusses the results related to the second question posed at the beginning of this thesis:

> What is the effect of ELMs on the dynamical behaviour of radiation fronts in the scrape-off layer?

Answering this question required a two-step approach. The code that existed in December 2019 did not include any sort of ELM functionality, and so the first goal was to build that functionality. First, this required determining the key features of ELMs that such a functionality should incorporate. Second, this required a practical understanding of the way that DIV1D implemented the physics that was discussed in chapter 2 in code form. The first of these requirements is elaborated on in section 5.1 of this chapter. Titled *Including ELM models in DIV1D,* it discusses basic parameters that were found to be key in describing ELMs. It then goes into two different ELM models, one of them from literature (Eich *et al*, 2017) and the other of devised as part of this thesis.

The self-built model was ultimately not employed in any of the simulations in this work—it was built as a smoother variant of the Eich *et al* version in the hope of reducing numerical noise but ultimately did not provide better results—but it is still an interesting model. It closely reproduced the ELM shape of Eich *et al*, but without discontinuities in the first time derivative, while retaining most of what made the other model nice to work with. It is retained in this work for those reasons. The second requirement, that of a practical understanding of DIV1D, is not discussed in this chapter as it is less a result and more a necessary prerequisite to implementing the ELM functionality. Some of it can be gleaned from appendix B, which describes the way that the code was implemented.

The second section, 5.2, is a case study for the MAST tokamak using this newly implemented ELM functionality. The reason for picking MAST is a rather pragmatic one. It is the device for which SD1D, the C++ code on which DIV1D is based, was originally developed, and consequently its parameters were used for many of the early simulations in DIV1D. It was a natural extension of this existing work likewise to take MAST as a proving grounds for the newly implemented ELMs. In order to subject the plasma to ELMs that were—if not realistic—at least physically plausible for MAST, back-of-the-envelope calculations were performed for an order-of-magnitude estimation of the relevant ELM parameters. These calculations are discussed in sub section 5.2.1. Using these back-of-the-envelope parameters as anchors several parameter scans were performed in both weakly detached (sub section 5.2.2) and fully detached plasmas (sub section 5.2.3). A key parameter in these calculations was the excess target fluence induced by the ELMs, which is an important indicator for target survival. Its magnitude was probed as a function of ELM amplitude (upstream fluence), repetition rate and upstream density.

5.1. Including ELM models in DIV1D

Edge-localised modes—which typically go by their abbreviation: ELMs—are a class of phenomena first observed in H-mode[9] tokamak discharges that periodically cause particle and heat losses from the core

[9] H-mode (high confinement mode) is a regime of magnetic confinement that was discovered in the early eighties (Wagner, 1982) in the ASDEX tokamak. It was and is a regime that differed from the contemporarily known mode of plasma operation, retroactively dubbed the L-mode (with the L for low), in that it had a higher core density

plasma. A review of ELMs and their effects in divertor tokamaks by Hill (1997) discusses some of their key features. There are three main types of ELMs—Types I, II and III—that are caused by different instabilities associated with the core plasma. They influence the heat and particle balances to a varying extent and they occur at different frequencies. The extent of their influence increases inversely with their occurrence rate. Large Type I ELMs can release as much as $10–15\%$ of the plasma thermal energy and density but they occur at a rate of $10\ \mathrm{Hz}$ or less. Small Type II and Type III ELMs, meanwhile, repeat at much higher frequencies in the $10^2–10^3\ \mathrm{Hz}$ range while their individual effects are too small to measure. ELM behaviour differs from device to device, and to determine how ELM behaviour extrapolates to bigger devices such as ITER scaling laws have been derived (Eich *et al*, 2017).

Since the core plasma is not included in DIV1D's domain, self-consistently modelling ELMs is out of the question. Furthermore, since ELMs are initiated somewhere in the core neither is it assumed that either ELM particles, momentum or energy suddenly appear along the divertor leg. Rather, we assume in this work that ELMs—just like the steady-state heat flux—present themselves as a time-dependent boundary condition on the upstream side of the domain. To impose such a time-dependent boundary condition, though, requires additional information. Particularly, it begs the question: "How do ELMs temporally evolve on the upstream side of the divertor leg?"

One model used in literature (Eich *et al*) to describe Type I ELMs at the target side (which is emphatically not the upstream) uses a simple temporal profile for the heat flux effected by an ELM. The ELM is described as a triangular waveform with two distinct phases: a linear rise over a duration of $\tau_{\mathrm{r(ise)}}$ followed by a linear fall over a duration $\tau_{\mathrm{f(all)}}$. Those timescales are coupled by further approximating that $\tau_{\mathrm{f}} = 2\tau_{\mathrm{r}}$. The area under the profile reflects the total fluence Q_{ELM} that impacts the divertor, i.e. the target heat flux integrated over the course of an entire ELM. This waveform takes the shape

$$q_{\mathrm{ELM}}(t) = \frac{2}{3}\frac{Q_{\mathrm{ELM}}}{\tau_{\mathrm{r}}} \times \begin{cases} \dfrac{t}{\tau_{\mathrm{r}}} & 0 \le t \le \tau_{\mathrm{r}} \\ 1 - \dfrac{t-\tau_{\mathrm{r}}}{2\tau_{\mathrm{r}}} & \tau_{\mathrm{r}} < t \le 3\tau_{\mathrm{r}} \end{cases} \tag{5.1}$$

and has the useful property that

$$\int_0^{3\tau_{\mathrm{r}}} q_{\mathrm{ELM}}(t)\mathrm{dt} = Q_{\mathrm{ELM}} \tag{5.2}$$

such that the rise time τ_{r} can be varied freely while keeping the total ELM fluence Q_{ELM} unchanged. The previously steady-state boundary condition for the X-point heat flux, $q_{\parallel,\mathrm{X}} = q_{\mathrm{s(teady)s(tate)}}$, can be amended using equation 5.1 to define a new and time-dependent boundary condition $q_{\parallel,\mathrm{X}}(t)$ as

$$q_{\parallel,\mathrm{X}}(t) = q_{\mathrm{ss}} + q_{\mathrm{ELM}}(t). \tag{5.3}$$

There is one oddity in the above line of reasoning. q_{ss} is defined at the X-point position, while $q_{\mathrm{ELM}}(t)$ was constructed solely based on target considerations. Separating either position is the full numerical

and a much flatter density profile in the core (the pedestal) that dropped off sharply near the edge plasma. H-mode is favourable to L-mode thanks to its comparatively high density and increased energy confinement time (which are two out of three factors in the triple product, equation 1.2). The trade-off, however, is that ELMs start popping up, which do not exist in L-mode. Despite ELMs, ITER is envisaged to run in H-mode (EUROfusion, 2019).

domain with a length of L that numbers in the tens (to potentially hundreds) of meters. One might argue that it would be better to use the *in situ* temporal evolution at the X-point rather than at the target to describe the upstream heat flux. This is a sentiment with which I would readily concur but, alas, cannot. The literature described above either discusses ELMs based on measurements at the target or based on global inventories. Putting that objection aside one might instead wonder what testable assumption would make the coupling of target heat flux profile to the upstream heat flux profile plausible. To this I propose the following. If the pulse as it is defined at the X-point persists up to the target while retaining its original shape, the target pulse shape gives a good indication for the upstream pulse shape in the context of DIV1D. This assumption appears to hold reasonably well. Figure 5.2 shows that, although the shape of the ELM becomes more erratic over the course of the flux tube, the general structure of a steep incline followed by a smoother decline is retained along the flux tube. This observation lends strength to the assumption that the pulse shape on the target side is representative for the pulse shape on the upstream side.

One feature of the triangular ELM model did not sit well, however. Those were the three discontinuities in the first derivatives with respect to time of the triangular model at $t = 0$, $t = \tau_\mathrm{r}$ and $t = 3\tau_\mathrm{r}$. They imply an instantaneous change to the rate at which the heat flux grows, which is unphysical. This unease was amplified by an early observation that the triangular ELMs left considerable oscillations in their wake in the solutions to all plasma parameters. They decayed over timescales that were an order of magnitude longer than the timescales of those ELMs. Although it later seemed that these oscillations appear to be a more universal feature of DIV1D, their timescales not tied to those of the ELMs at all, contemporary understanding motivated the construction of a discontinuity-free version of the model by Eich *et al.* Although the model was not actually used for the simulations described in the next section, it has been implemented in DIV1D and exists in the current code version, so a description of it is included here.

As such models can be constructed in a plethora of ways, the mathematical architect is free if not encouraged to be somewhat picky in defining its properties. There are three must-haves. Firstly, it should be at least once differentiable with respect to time. Secondly, it should start out from zero and go to zero. Thirdly, its profile should be quantitatively similar to that described in equation 5.1. Additionally, we are free to define as many nice-to-haves as we see fit. For one, to limit complexity it would be nice not to have to resort to a piecewise formulation. Likewise, it would be nice if the model were analytically normalisable. Finally, it would be nice to limit the number of independent parameters to those with a clear physical interpretation. Concretely, this translates to having just two parameters: one to determine the ELM pulse length and one to determine the ELM amplitude. Both the must-haves and all but one of the nice-to-haves were defined *a priori*[10], and they combinedly motivated the choice of the Maxwell-Boltzmann distribution,

$$f(x_i) = \sqrt{\frac{2}{\pi}}\frac{x_i^2}{a^3}\exp -\frac{x_i^2}{2a^2}. \tag{5.4}$$

Out of the box it checks two out of three must-haves. It is infinitely differentiable with respect to an arbitrary argument x_i, and it starts out zero and asymptotically returns there. Additionally, it checks all the nice-to-haves. It is non-piecewise, analytically normalised by definition, and its width is determined

[10] The statement that it preferably not be piecewise was incentivised later by a rather unsuccessful attempt to render equation 4.1 continuous through the addition of three parabolas to bridge the discontinuities.

by a single parameter a while its normalised nature allows it to be multiplied by any amplitude. A $5/6$ score would have been acceptable if not for the fact that one of the must-haves was missing: the Maxwell-Boltzmann distribution is not sufficiently asymmetric to resemble equation 5.1. Because it fit so well within all other design constraints it was not immediately discarded, though. And if one distribution does not cut it, a couple just might. Rather than using a single Maxwell-Boltzmann distribution, then, a summation over Maxwell-Boltzmann distributions was defined as

$$q_{\mathrm{ELM}}(t) = Q_{\mathrm{ELM}} \frac{1}{N} \sqrt{\frac{2}{\pi}} \sum_{n=1}^{N} \frac{t^2}{a_n^3} \exp -\frac{t^2}{2a_n^2}. \tag{5.5}$$

where the widths of consecutive-order terms are determined by

$$a_n = \Delta^{n-1} \sqrt{\frac{\tau_{\mathrm{r}}^2}{2}}. \tag{5.6}$$

Note that the $n = 1$ term reduces to $a_n\sqrt{2} = \tau_{\mathrm{r}}$, thus setting the time of maximum heat flux to τ_{r} for that term (and for the entirety of equation 5.5 in case $N = 1$). Δ is a new shaping parameter that dictates a constant spacing between consecutive maxima, and thus allows one to tune the level of asymmetry. For $\Delta = 1$ the original Maxwell-Boltzmann distribution is recovered regardless of the number of terms N. Higher-order terms move the time of maximum heat flux to a somewhat later point assuming $\Delta > 1$. For brevity's sake and to avoid any semblance of a connection to an underlying energy distribution in some form or the other, this model has simply been dubbed the gaussian model. Both the triangular model and the gaussian model have been plotted in figure 5.1 for the same choices of τ_{r} and Q_{ELM}. They agree very nicely.

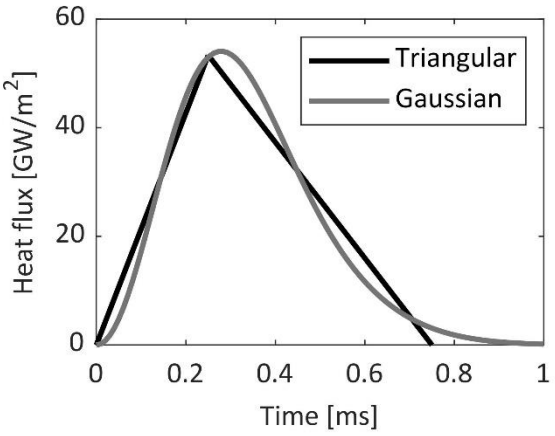

Figure 5.1. A comparison of two models for an ELM's heat flux in a high-power ITER-like scenario ($\tau_{\mathrm{r}} = 250$ µs and $Q_{\mathrm{ELM}} = 20\ \mathrm{MJ/m^2}$ (Eich *et al*, 2017)). The triangular model, shown in black, was obtained from Eich *et al* (2017) while the gaussian ($N = 2, \Delta = 1.4$) version derives from equation 5.5 in the present work.

One might point out that the introduction of the above summation invalidates one nice-to-have: namely, that the independent parameters should number two. This is partly true. However, N and Δ need only be tuned once. They additionally allow some freedom in adapting the ELM to be more or less asymmetric. This has proved helpful once already, as an earlier version of the ELM model was tuned to take after TCV bolometry measurements (Veres *et al*, 2009) rather than after the triangular model by Eich *et al*. By setting shaping parameters of $N = 4$ and $\Delta = 1.9$ it was possible to reproduce the much

steeper rise and longer tail (stronger asymmetry) measured in TCV during ELMs as readily as it was to reproduce the triangular shape. When it later turned out that the underlying assumption—that the global evolution of the plasma radiation was reflected in the local behaviour of the divertor—was not valid for that TCV measurement as it had included significant ELM-induced losses directly to the vessel wall, the model was tuned to the present parameters without much trouble. The versatility that they allow speak in favour of these shaping parameters.

Figure 5.2 compares two simulations for a fully detached plasma ($n_{\mathrm{u}} = 5.0 \times 10^{19}\ \mathrm{m}^{-3}$, $L = 20\ \mathrm{m}$ and $q_{\parallel,\mathrm{X}}(0) = q_{\mathrm{SS}} = 50\ \mathrm{MW\ m}^{-2}$), one of which was exposed to a triangular ELM and the other one of which to a gaussian ELM. Time traces of the heat flux are given for both simulations at five different locations along the divertor leg, including the upstream position. As the heat flux progressed through the divertor leg the time traces became increasingly more erratic. Outside of expectations, the simulation that used the triangular ELM actually appeared less noisy than its gaussian counterpart. Of course, figure 5.2 is only one example but having tried both types of ELMs for a variety of shots, the same behaviour was recovered. While the smoother ELM variant was developed with the explicit goal of reducing instabilities during the simulation, it did not end up doing much to alleviate matters in the case of DIV1D. It displayed feature parity at best, but often performed worse than its simple triangular counterpart. As a result, the ELM simulations performed in the rest of this chapter were all performed using the triangular variant by Eich *et al*.

5.2. A case-study: MAST like ELMs in DIV1D

5.2.1. Back-of-the-envelope parameter estimates

As a case study for simulating ELMs, the MAST tokamak was taken into consideration. ELMs have been thoroughly characterised in MAST (e.g. Kirk *et al*, 2004), and that information can be used to determine both the ELM duration, its amplitude and its repetition rate used by DIV1D. Since measurements of ELM power seem to be predominantly carried out at the divertor target, this brings one issue to the fore. ELMs in DIV1D are defined on the upstream side of the plasma. In the previous section it was seen that the general pulse shape of the ELM remained more or less the same when moving from upstream to target in the context of DIV1D. The pulse shape for the ELM defined at the target by Eich *et al* (2017) and the smoother variant of this thesis, then, could be expected to remain the same from upstream to downstream. However, the same does not necessarily go for the ELM amplitude. Below a simplistic calculation is performed to calculate this amplitude. The goal of this calculation is to find a plausible number for the ELM amplitude—not a particularly accurate one—to get a sense of the contours of the physically plausible parameter space that DIV1D could probe.

Table 5.1. Plasma parameters for a typical MAST discharge in L-mode (Meyer, 2016).

Minor radius a	0.65 m
Major radius R	0.85 m
Core density n	$3 \times 10^{19}\ \mathrm{m}^{-3}$
Core ion temperature T_{i}	3 keV (35 MK)
Core electron temperature T_{e}	2 keV (23 MK)
Plasma current I_{p}	1.2 MA

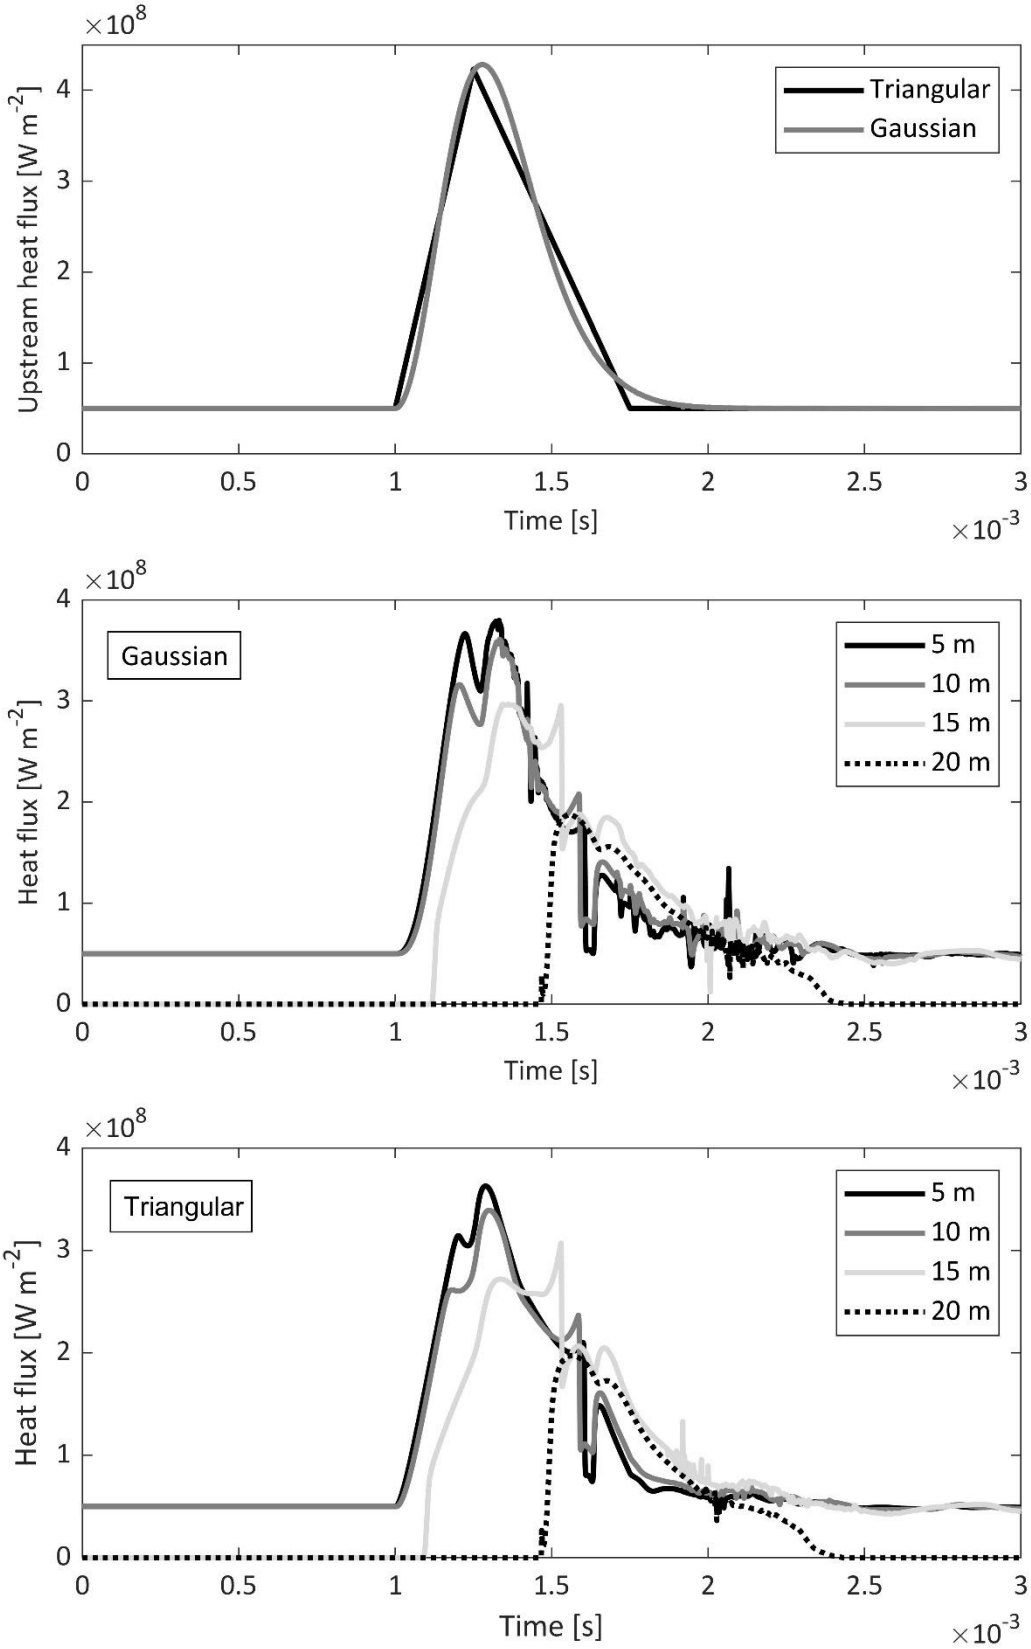

Figure 5.2. A comparison of two DIV1D simulations for fully detached plasmas ($n_\mathrm{u} = 5.0 \times 10^{19}\ \mathrm{m}^{-3}$, $L = 20\ \mathrm{m}$ and $q_{\parallel,\mathrm{X}}(0) = q_\mathrm{SS} = 50\ \mathrm{MW\ m}^{-2}$) subjected to like ELMs ($Q_\mathrm{ELM} = 140\ \mathrm{kJ\ m}^{-2}$ and $\tau_\mathrm{r} = 250\ \mu\mathrm{s}$) with two different shapes. The top graph shows both ELMs on the upstream side of the domain. The second and third graphs show how both ELMs propagate through the flux tube, giving time traces at $5\ \mathrm{m}$, $10\ \mathrm{m}$ and $15\ \mathrm{m}$ beyond the upstream point, and at the target itself ($20\ \mathrm{m}$). The time traces depict similar behaviour, although the triangular ELM seems less heavily affected by fluctuations. In either situation, the time traces grow more erratic as the heat flux progresses through the flux tube. The general shapes, however,

seem to be retained. Although the heat flux for the $15\ m$ position does not revert to its pre-ELM position over the time plotted here, it does eventually return to zero but this takes time as the detachment front has to come all the way back from the target.

In order to determine the upstream heat flux that goes into the flux tube, several pieces of information can be combined. Kirk *et al* (2004) note that a big ELM in MAST can carry from the core as much as three percent of the total thermal energy in the plasma. However, they do not quantify that total thermal energy E_{th}. By taking the plasma parameters of MAST and typical values for the core density and temperatures (table 5.1 (Meyer, 2016)), one could approximate the thermal energy density as

$$E_{\mathrm{th}}/V = \frac{3}{2} n k_{\mathrm{B}} (T_{\mathrm{e}}[\mathrm{K}] + T_{\mathrm{i}}[\mathrm{K}]) = \frac{3}{2} n e (T_{\mathrm{e}}[\mathrm{eV}] + T_{\mathrm{i}}[\mathrm{eV}]) \tag{5.7}$$

where the volume V follows from a cylindrical approximation of the tokamak volume, with a and R the minor and major radii, as

$$V = \pi a^2 \times 2\pi R. \tag{5.8}$$

This overestimates the thermal content somewhat because the temperature and densities drop outside of the core. However, the parameters in table 5.1 are for L-mode plasmas, which have lower densities than the H-modes plasma which produce ELMs. Opposite to the other effect, this causes us to underestimate the thermal content. For these parameters it follows that $E_{\mathrm{ELM}} = 3\% \times E_{\mathrm{th}} = 7.8\ \mathrm{kJ}$. To move from E_{ELM} to q_{ELM}, however, additionally requires both geometrical and temporal considerations. With respect to the latter, an ELM rise time of $250\ \mu\mathrm{s}$ is assumed in keeping with what is expected for ITER (Eich *et al*, 2017). This translates to a total pulse time of $\tau_{\mathrm{ELM}} = 750\ \mu\mathrm{s}$ via the ELM prescribed in section 5.1, and thus to an average extra power load through the SOL of $P_{\mathrm{ELM}} = E_{\mathrm{ELM}}/\tau_{\mathrm{ELM}} = 10.4\ \mathrm{MW}$. The area of the SOL depends on the toroidal length of MAST, which is $2\pi R$, and on the radial width of the SOL, which is a highly non-trivial parameter. For MAST, a scaling law has been established that gives the radial width $\lambda_{\mathrm{q}}\ [\mathrm{mm}]$ of the scrape-off layer as a function of plasma current $I_{\mathrm{p}}[\mathrm{MA}]$ and SOL power $P_{\mathrm{SOL}}[\mathrm{MW}]$ (Thornton, Kirk and the MAST Team, 2014) as

$$\lambda_{\mathrm{q}} = 4.57(\pm 0.54) I_{\mathrm{p}}^{-0.64(\pm 0.15)} P_{\mathrm{SOL}}^{0.22(\pm 0.08)}. \tag{5.9}$$

Although this scaling has been determined for the inter-ELM regime one might nonetheless try to extend it to the ELM regime under the assumption $P_{\mathrm{SOL}} \approx P_{\mathrm{ELM}}$ which yields a SOL width of $6.8\ \mathrm{mm}$ and thus an effective area of $0.036\ \mathrm{m}^2$. The total energy fluence per square meter, finally, is $2.15 \times 10^5\ \mathrm{J/m^2}$ and the average extra heat flux is $287\ \mathrm{MW/m^2}$. A different method using the ELMy H-mode scaling law in ITER IPB98(y,2) (ITER Physics Expert Group on Confinement and Transport *et al*, 1999) yields a slightly lower total loss of 2.2% of the thermal content for the biggest of ELMs (Kirk *et al*, 2004). This number is well within the same ballpark. This results in an average extra energy fluence and heat flux that differ by a factor of $(2.2/3)^{0.78} \approx 0.8$. The less-than-one dependence on the thermal plasma results from the fact that the radial width also weakly scales with the power through the SOL as described by equation 5.9. This would bring the ELM heat flux down to about $225\ \mathrm{MW/m^2}$ and the energy fluence to $1.7 \times 10^4\ \mathrm{J/m^2}$. Kirk *et al* (2004) also note that, for the largest ELMs, the ratio between the ELM and inter-ELM power loads at the divertor strike point is greater than five. Dudson *et al* (2019) state that a typical parameter for the upstream heat load in MAST is about $50\ \mathrm{MW/m^2}$, which means that the ratios of ELM to inter-ELM for the above calculations, $(50 + 287)/50 = 6.7$ and $(50 + 225)/50 = 5.5$, appear to reasonably confirm the experiments described by Kirk *et al* (2004).

5.2.2. Weakly detached plasmas subjected to single-pulse ELMs

Using these values, we ran a series of simulations for the same two scenarios that have appeared earlier in this thesis. One of them featured a weakly detached plasma and the other a fully detached one. In either situation a previously obtained steady-state solution was used as the initial solution. We set an extra 1-ms delay to make sure that the run really was steady-state, and afterwards pulsed a single triangular ELM on the upstream side. The ELM duration was fixed at $750\ \mu\text{s}$ while the upstream fluence (Q_{ELM} in equation 5.5) was varied between $10\ \text{kJ/m}^2$ and $140\ \text{kJ/m}^2$. Using an upstream fluence of $200\ \text{kJ/m}^2$ was also attempted but the simulation would not converge in that case. Results for the weakly detached scenario ($n_{\text{u}} = 3.9 \times 10^{19}\ \text{m}^{-3}$, $L = 20\ \text{m}$, $q_{\text{X,}\parallel}(0) = q_{\text{ss}} = 50\ \text{MW m}^{-2}$) have been plotted in figure 5.3.

Additionally, the line $y = x$ has been plotted to show the hypothetical scenario in which the full excess upstream fluence ends up at the divertor target in figure 5.3. This is an interesting reference point because all data points that live above it indicate an ELM target fluence in excess of the amount that was added on the upstream. Although this might sound strange, it is actually not. As the plasma fully reattaches the temperature rises and higher temperatures relate to less effective sinks. Part of the steady-state heat flux which was originally removed by those sinks thus impinges on the divertor target in addition to the extra heat flux due to the ELM. The combination of either effect can result in fractional ELM fluences at the target above one. This is one of the reasons why reattachment is so dangerous.

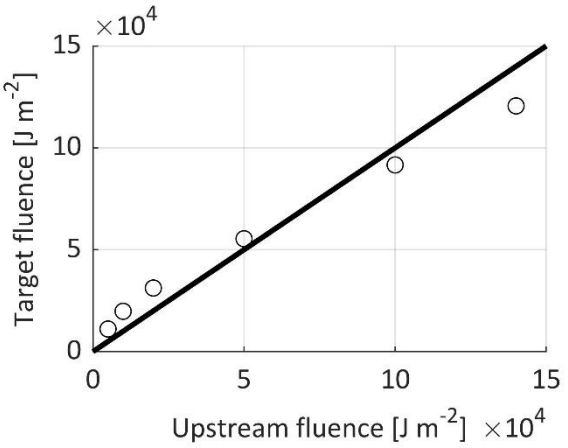

Figure 5.3. The extra target fluence, i.e. the extra time-integrated heat flux, due to a single ELM as a function of the upstream ELM fluence for a weakly detached plasma ($n_{\text{u}} = 3.9 \times 10^{19}\ \text{m}^{-3}$, $L = 20\ \text{m}$, $q_{\text{X,}\parallel}(0) = q_{\text{ss}} = 50\ \text{MW m}^{-2}$). Also drawn is the line $y = x$, which serves as a reference. It indicates the (hypothetical) scenario in which the ELM-induced heat flux on the upstream side propagates unabatingly to the target.

At higher upstream fluences, the opposite effect is observed: data points live below the diagonal, reflecting that not all of the extra inputted heat flux reaches the divertor target. This appears to be the result of the plasma fully detaching in the wake of the ELM, which can be seen to happen in figure 5.4. It shows time traces of three parameters: the heat flux at the target, the temperature at the target and the position of the detachment front. After having run the simulation for an initial $1\ \text{ms}$ we initiated an ELM on the upstream side of the domain. This is indicated by the narrow vertical black lines in the plots. The response of the target appears instantaneous on the timescale plotted in figure 5.4. (and it almost is) but there is a brief $15\text{-}\mu\text{s}$ delay before the heat flux at the target started to pick up. A large initial peak was observed in both the heat flux and the temperature at the target.

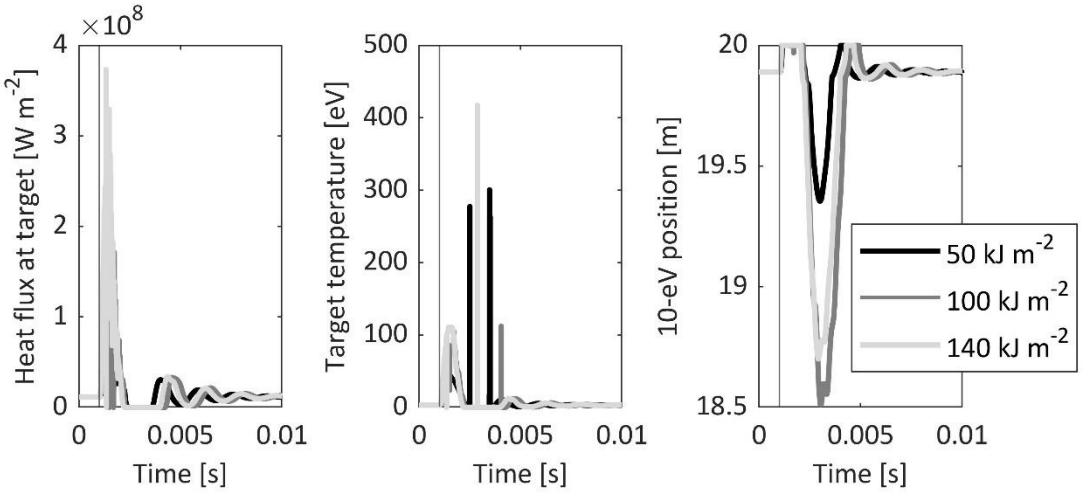

Figure 5.4. The target heat flux, target temperature and position of the detachment front (10-eV position) as functions of time at the highest three upstream ELM fluences in figure 5.3. Also plotted are vertical lines that indicate the starting time of the ELM. What catches the eye is that the plasma detached from the divertor target in the wake of the ELM's initial perturbation for all three solutions. The vertical peaks in the target temperature during the detached period were often observed in DIV1D and appear to be numerical artefacts; they did not significantly affect the solution (since the plasma density is almost zero at the target in these situations) and usually disappeared after one time step.

In case of the 140-kJ-m^{-2} ELM plotted in light grey, the heat flux peaked at a value in excess of $350\ \text{MW m}^{-2}$, which is more than seven times the inter-ELM value, while the target temperature peaked at just over $100\ \text{eV}$. During this time, the front had fully reattached which is reflected in the 10-eV position having moved to the target. Strictly speaking, once at the target, we can no longer call it the 10-eV position as the temperature is higher along the entire domain. The parameter reflects the position of the minimum temperature in these cases. After this initial peak had passed, the target temperature took a dive and reached (numerical) zero which means that the plasma—first only weakly detached, then reattached—had entered a fully detached state with the detachment front receded away from the target. As a result, no heat reached the divertor target at all over a period of $1\text{–}2\ \text{ms}$. This period was shorter for a 50-kJ-m^{-2} than for a 100-kJ-m^{-2} or 140-kJ-m^{-2} ELM, but it does not seem to be the ELM's intensity alone that dictates the extent and duration of detachment as the 140-kJ-m^{-2} ELM resulted in a shorter detachment period than the 100-kJ-m^{-2} ELM. There does seem to be a relation between the extent of detachment (quantified by how far the 10-eV position moved upstream) and the duration of detachment, with a further receded detachment front relating to a longer period of detachment. This relation is probed more explicitly later in this section.

The transition from the above-diagonal to below-diagonal behaviour in figure 5.3 is tentatively linked to whether or not the target temperature dips to $0.1\ \text{eV}$ or not over the course of the simulation. In figure 5.5 the same parameters are plotted for the first three data points of figure 5.3. What can be observed here is that the size of the initial power perturbation by the ELM appears to dictate the size of the fluctuations that are left in its wake. Moving from the 5-kJ-m^{-2} through the 10-kJ-m^{-2} to the 20-kJ-m^{-2} ELM, we saw the fluctuation amplitude increase. As this fluctuation amplitude increased, the global minimum of the target temperature with respect to time—found in the first period of the oscillation in the wake of the ELM—crept ever nearer to $0.1\ \text{eV}$. For the last data point above the line in figure 5.3, this minimum was still marginally above zero. Comparing figures 5.4 and 5.5, we see a last thing stand out. While the detachment front location oscillated rather smoothly in the simulations of figure 5.5, it behaved much more erratically in those of figure 5.4. While the detachment fronts initially slowed as they approached the minimum temperature in all three cases, they suddenly sped up as temperatures reached $0.1\ \text{eV}$.

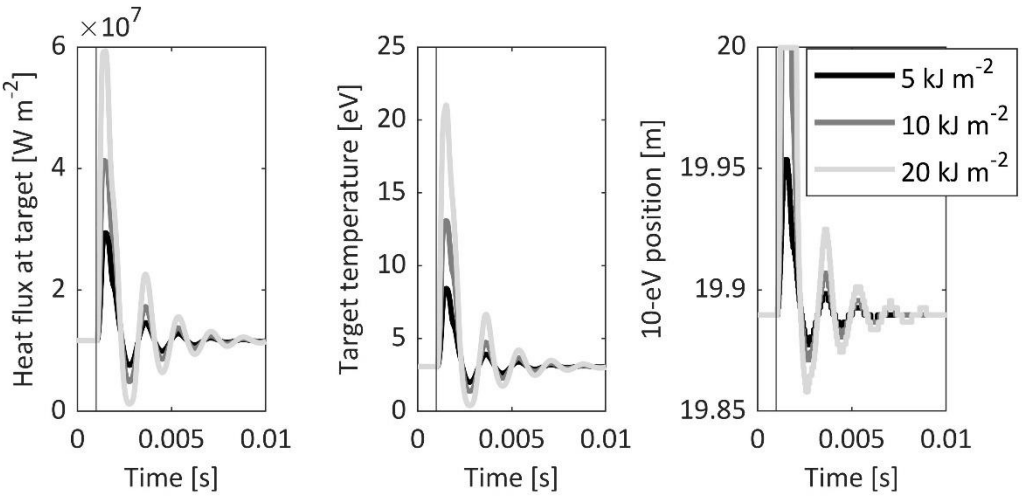

Figure 5.5. The target heat flux, target temperature and position of the detachment front (10-eV position) as functions of time at the lowest three upstream ELM fluences in figure 5.3. Also plotted are vertical lines that indicate the starting time of the ELM. Unlike in the previous picture, the plasmas did not fully detach in the wake of the ELM's initial perturbation. As the ELM intensities grew, the fluctuations left in their wake did grow with the biggest ELM shown here (with a fluence of $20\ \mathrm{kJ\ m^{-2}}$) having a target temperature only marginally above zero.

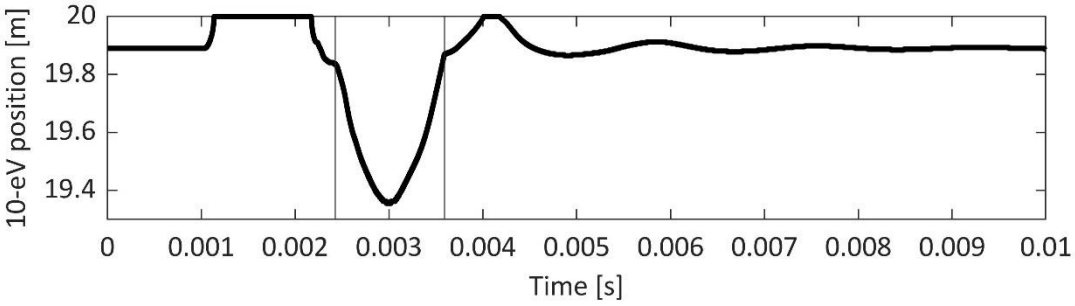

Figure 5.6. A zoom-in on the detachment front (10-eV position) as a function of time for the 50-kJ-m^{-2} ELM. Also plotted are two vertical lines that indicate the start and end of the time domain in which the target temperature had dropped to zero. These lines coincide with a pair of kinks in the solution—a strong localised change in the second time derivative interrupts the flattening of the curve—that reflect that the detachment front suddenly accelerates away from the target when the target temperature reaches zero.

To better illustrate this, the position of the detachment front of the 50-kJ-m^{-2} ELM has been replotted in figure 5.6. Also plotted are two vertical lines that indicate the boundaries of the time domain in which the temperature at the target was $0.1\ \mathrm{eV}$, which is the period in which the plasma was fully detached. They coincided with a sudden changes in the second time derivative. Rather than slow down further, the front suddenly picked up speed and accelerated away from the target, which was in breaking with the observed trend prior to the temperature dropping to $0.1\ \mathrm{eV}$. This indicates that something is going on as the temperature drops to $0.1\ \mathrm{eV}$.

5.2.3. Fully detached plasmas subjected to single-pulse and recurring ELMs

A similar parameter scan as in figure 5.3 was performed for a fully detached plasma ($n_\mathrm{u} = 3.9 \times 10^{19}\ \mathrm{m^{-3}}$, $L = 20\ \mathrm{m}$, $q_{\mathrm{X},\parallel}(0) = q_\mathrm{ss} = 50\ \mathrm{MW\ m^{-2}}$). Zero target fluence is observed at the lower ELM intensities up to $20\ \mathrm{kJ\ m^{-2}}$ while the slope between the last two data points is somewhat under unity in keeping with what was observed in figure 5.3. This suggest that a threshold value exists under which the detachment front provides a sufficient buffer to incoming ELMs.

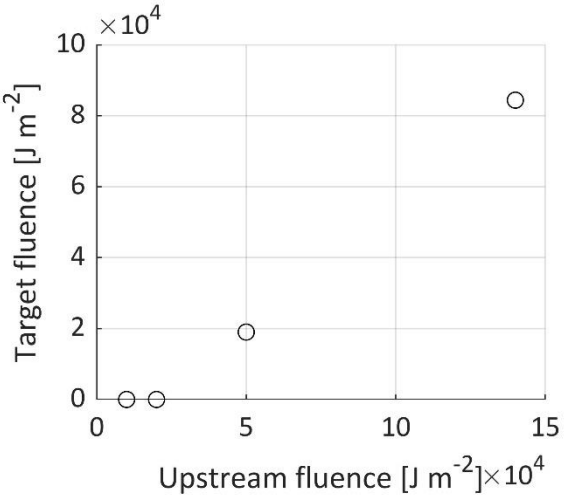

Figure 5.7. The extra target fluence due to a single ELM as a function of the upstream ELM fluence for a fully detached plasma ($n_{\mathrm{u}} = 5.0 \times 10^{19}\ \mathrm{m}^{-3}$, $L = 20\ \mathrm{m}$, $q_{\mathrm{X,\parallel}}(0) = q_{\mathrm{ss}} = 50\ \mathrm{MW\ m}^{-2}$). The target fluence remains zero initially. This reflects that the ELM is not strong enough to cause reattachment of the plasma. For stronger ELMs, the plasma does reattach as exemplified by the non-zero target fluences in those cases.

To probe this effect as a function of the upstream density, we selected part of an existing density scan by Westerhof (unpublished) and subjected it to a fixed-fluence ELM of $140\ \mathrm{kJ\ m}^{-2}$. The results are plotted in figure 5.8. There is a strong linear tendency to the relationship between the target fluence and the upstream density. As the upstream density was increased while the upstream ELM fluence was kept constant, the target fluence was seen to decrease from around $140\ \mathrm{kJ\ m}^{-2}$, i.e. the full starting fluence, at a density of $4.5 \times 10^{19}\ \mathrm{m}^{-3}$ to around $0\ \mathrm{kJ\ m}^{-2}$ at a density of $5.7 \times 10^{19}\ \mathrm{m}^{-3}$. This effect appeared strongly correlated with the position of the detachment front, which is plotted in figure 5.9 over a slightly narrower range. The figure suggests a likewise linear correspondence between the upstream density and the position of the detachment front. As both the target fluence and the position of the detachment front appear to be almost linearly related to the plasma density, this leads one to conclude that the target fluence also depends linearly on the position of the detachment front for a fixed ELM size.

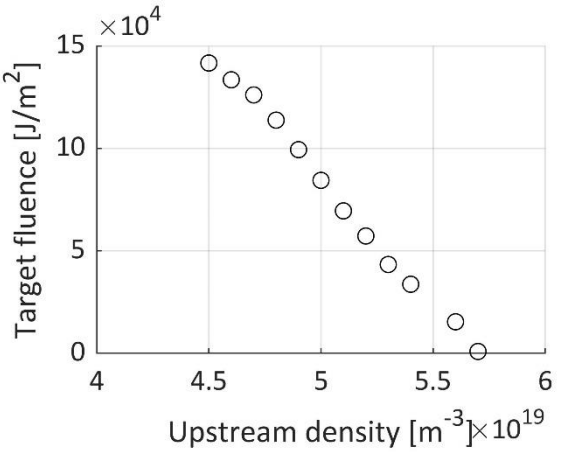

Figure 5.8. Target fluence due to a 140-kJ-m^{-3} ELM as a function of upstream density. A range of previously obtained steady-state solutions was by Westerhof (unpublished) was perturbed with this fixed-amplitude ELM. At a density of $4.5 \times 19^{19}\ \mathrm{m}^{-3}$ the full ELM-induced upstream fluence ends up at the target. At a density of $5.7 \times 10^{19}\ \mathrm{m}^{-3}$ this value has dropped to zero.

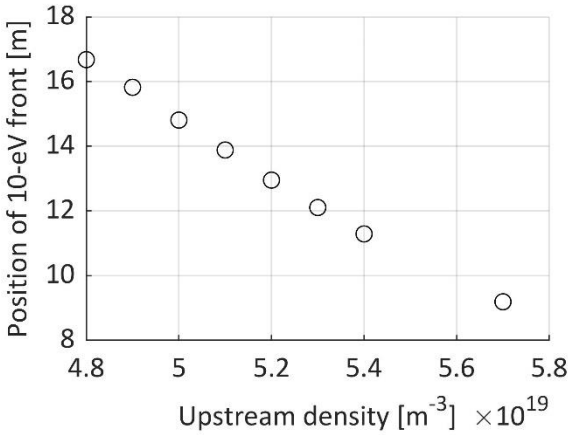

Figure 5.9. Position of the 10-eV front, as a proxy for the detachment front, as a function of the upstream density for a part of the data point in figure 5.8. The divertor target is situated at $20\ \mathrm{m}$. There is an almost linear correspondence between the detachment front's position and the upstream density.

ELMs are recurring phenomena, however, which makes one wonder when it is realistic to draw conclusions based on single-pulse ELMs. Judging by figure 5.4, even the strongest ELM at $140\ \mathrm{kJ\ m^{-2}}$—which is close to the back-of-the-envelope value found in section 5.2.1—do not perturb the system much beyond $10\ \mathrm{ms}$, which might motivate one to conclude that ELMs with a repetition rate of less than $100\ \mathrm{Hz}$ can adequately be described by single-pulse ELMs as the effects of the previous ELM will have died out by the time that a next ELM presents itself. As the large Type I ELMs have a repetition rate well below this value (Hill, 1997), those can adequately be described as single pulses. However, as one moves into Type II and III ELMs, the repetition rate tends to increase to values well above $100\ \mathrm{Hz}$. In those cases, the steady-state heat flux of $50\ \mathrm{MW\ m^{-2}}$ might be underestimating the true situation. We hypothesised, rather, that a quasi-steady state would be obtained in the regime in which consecutive ELMs did not overlap on the upstream side, i.e. had a repetition rate of less than $1/750\ \mu\mathrm{s} = 1.3\ \mathrm{kHz}$, but whose effects over the length of the flux tube had not yet died out. It was this regime between 10^2–$10^3\ \mathrm{Hz}$ that would potentially display such quasi–steady-state behaviour.

To this end an additional option was implemented in DIV1D which allows the user to switch between a single-pulse or recurring ELM. If the latter is chosen, the user is additionally requested to input how many time steps should be between consecutive ELMs. This details of this implementation are included in appendix B and are included in the current master branch[11] of the DIV1D code. Putting this newly implemented feature to use, a scan was performed to determine the effect of the ELM frequency on the target fluence for a fully detached plasma. Frequencies ranging from $250\ \mathrm{Hz}$ to $1000\ \mathrm{Hz}$ were chosen at a fixed upstream ELM fluence of $20\ \mathrm{kJ\ m^{-2}}$. The results of this scan have been plotted in figure 5.10. Rather than plotting the absolute fluence at the target, we instead plotted the target fluence in units of the upstream fluence and called it the fractional target fluence. This was necessary as the absolute fluence trivially rises with the number of ELMs per second, which says nothing about the percentage of the upstream fluence that actually ends up at the target as a function of the repetition rate. The fractional target fluence corrects for this effect.

[11] The term "master branch" is used in a git-based version control workflow typically to denote the default and stable version of a code. Although DIV1D versions were originally controlled by SVN, we decided to port the entire repository to git during the summer as distributed version control (Git) has some marked advantages over centralised version control (SVN) such as the ability to make branches (prototype code versions) locally rather than on the server side and the option to use intuitive graphical user interfaces like Gitlab, which was recently deployed at DIFFER.

What is observed in figure 5.10 is reminiscent of figure 5.7. At ELM frequencies up to $400\ \mathrm{Hz}$ the fractional target fluence is zero, indicating that the plasma remains fully detached despite recurring perturbations to the heat flux. From $500\ \mathrm{Hz}$ onwards the fractional target fluence starts rising, which indicates the onset of reattachment. At a frequency of $1000\ \mathrm{Hz}$, it sits at a value just shy of 0.9, meaning that most of the ELM-induced heat flux reaches the target unabatedly in that case. Also plotted are equivalent increases of the upstream heat flux equal to the time-averaged ELM fluences for the three non-zero data points. The behaviour is not exactly equal—particularly not for the $500\ \mathrm{Hz}$ simulation—but they do grow closer together as the ELM frequency increases. This seems to corroborate the suggestion that some sort of quasi-steady state is reached.

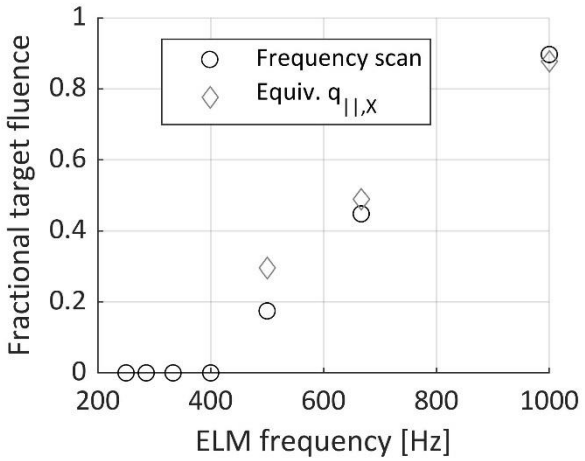

Figure 5.10. The fractional target fluence as a function of the rate at which ELMs recur, the ELM frequency, for a fully detached plasma ($n_\mathrm{u} = 5.0 \times 10^{19}\ \mathrm{m}^{-3}$, $L = 20\ \mathrm{m}$, $q_{\mathrm{X},\parallel}(0) = q_\mathrm{ss} = 50\ \mathrm{MW\ m}^{-2}$) at a fixed upstream ELM fluence of $2 \times 10^4\ \mathrm{J\ m}^{-2}$ with an ELM duration of $0.75\ \mathrm{ms}$ is plotted as the black circles. The absolute target fluence is not as indicative here as it depends linearly on the ELM frequency and will thus vary trivially as the system is subjected to more and more ELMs. The fractional target fluence, however, is calculated as the target fluence over the upstream fluence and thus impervious to that effect. Still, at higher ELM frequencies the fractional target fluence is seen to increase. ELM frequencies are low enough ($< 1/0.75\ \mathrm{ms} = 1.3\ \mathrm{kHz}$) over the entire range that consecutive ELMs never overlap on the upstream side. Also plotted are some ELM-less simulations for which the time-averaged extra ELM fluence was added as an increased upstream heat flux. These simulation use the same parameters as above (but no ELMs) and upstream heat fluxes of $q_\mathrm{ss} = 60\ \mathrm{MW\ m}^{-2}$, $q_\mathrm{ss} = 63.3\ \mathrm{MW\ m}^{-2}$ and $q_\mathrm{ss} = 70\ \mathrm{MW\ m}^{-2}$ for $500\ \mathrm{Hz}$, $667\ \mathrm{Hz}$ and $1000\ \mathrm{Hz}$ respectively. In these cases the fractional target fluence was calculated as the target heat flux at the end of the simulation over the added heat flux on the upstream with respect to the base scenario (so $+10\ \mathrm{MW\ m}^{-2}$, $+13.3\ \mathrm{MW\ m}^{-2}$ and $+20\ \mathrm{MW\ m}^{-2}$).

6. Summary and discussion

So what was the sense in doing all these calculations and plotting all these graphs? The title of this thesis is *Statics and dynamics of detachment in 1D*. How do the results of this thesis relate to that? Three chapters of results dealt with numerical considerations (chapter 3), steady-state solutions (chapter 4) and ELM dynamics. This chapter discusses the key findings of these chapters and their implications.

6.1. Code verification

A chapter on *Code verification* was included in this thesis to illustrate the most important steps that were taken to obtain confidence in the results of later chapters. There are several ways to verify that a model does what it was designed to do.[12] According to Thacker *et al* (2004) the most important step in what they call *calculation verification* is to perform a grid or time convergence study. Such a test checks whether a code yields increasingly accurate results for increasingly finer discretisation.

Section 3.1 discussed such a grid convergence scan. The goal was to understand to what extent the physical solutions were influenced by the degree of discretisation. It showed that target parameters (temperature and density) changed significantly for lower degrees of discretisation up to several 100s of grid points, while the upstream temperature stayed almost constant. The rate of change in the target parameters decreased with increasing discretisation. By 3500 grid points the target parameters were mostly impervious to further changes of the degree of discretisation. As one would expect the discretised equations to get ever nearer to their continuous counterparts as the degree of discretisation increased, it is a good sign that changes became slimmer at higher degrees of discretisation. Putting numbers to these observations, it was found that the target parameters changed by a factor of two from 100 to 3500 grid points, while they only varied by 10% from 500 to 3500 grid points and as little as 5% from 1000 to 3500 grid points. The grid convergence scan thus gives no cause for worry.

As an additional step to the code verification we looked at the effect of artificial viscosity—both for detached and attached cases—which is somewhat related to the above discussion of the grid convergence. It was formulated as a diffusion in index space (Dudson, 2016). One cannot help but see the similarity between equation 3.1 and a second-order central difference scheme of a function $f(x)$,

$$\frac{\mathrm{d}^2}{\mathrm{d}x^2} f(x) \approx \frac{f(x+h) - 2f(x) + f(x-h)}{h^2} \tag{6.1}$$

with h a finite difference. However, equation 3.1 does not include h^2 in the denominator and that makes it crucially different. As h in equation 6.1 becomes smaller the difference between $f(x)$ and its neighbours becomes ever smaller and thus the numerator shrinks. However, this is offset by the denominator. Removing that denominator as in equation 3.1, the effect of the artificial viscosity term becomes increasingly smaller for larger grids (which have smaller h). One of the drivers of the changes observed in the grid convergence scan, then, might be the artificial viscosity term. This observation is strengthened by two simulation findings. One of these was the observation in figures 3.2 and 3.4 that

[12] Note that a verified code does not constitute a physical code *per se*. One could purposefully design a code that yields unphysical results which nonetheless passes the steps of *verification* as it does what it was designed to do. Only during the *validation* step—in which the model is tested against reality—do these codes fail.

the momentum balance was much more strongly affected by a changing grid than the power balance. Second was the observation that increasing the artificial viscosity caused an opposite reaction of the target parameters compared to the reaction incurred by increasing the degree of discretisation. For an artificial viscosity in the range $\eta_{\mathrm{art}} = 0\text{–}1$, which is the range in which the simulations presented in this thesis have been carried out, a difference in the target parameters of maximally 3% was observed.

Finally, the sensitivity of the detachment front to changes in the artificial viscosity was also studied. It showed that no significant changes to the detachment front were occurred in the range between $\eta_{\mathrm{art}} = 0\text{–}5$ but that for higher values the changes could become considerable. Initially unexpected, it also showed considerable changes to the upstream temperature—which had been unchanging during the grid convergence scan and the artificial viscosity scan for an attached plasma. It was eventually found to be a consequence of the movement of the detachment front—and the ensuing change in the connection length between X-point and detachment front—rather than a numerical aberration. All in all, there is no reason to suggest that the artificial viscosity is a serious cause for worry at the values employed for the simulations in chapters 4 and 5.

6.2. Radial losses

Chapter 4 narrated the steps that were taken to compare DIV1D with SOLPS-ITER simulations by Wensing (2020), and what consequent changes where implemented to DIV1D when that comparison turned out to be lacking. In the initial comparison, we took 2PM independent parameters—upstream density n_{u}, upstream heat flux $q_{\parallel,\mathrm{X}}$ and connection length L—as well as some other parameters like the carbon concentration from a SOLPS-ITER simulation and fed them to DIV1D. Then, we compared the outputs of either code to see how well they agreed on the 2PM dependent parameters—the upstream and target temperatures T_{u} and T_{t}, and the target density n_{t}. Despite prolonged tweaking to all and any possible DIV1D parameters, they could not be made to agree. Something needed to be changed to the code, that much was clear. This something, initially, was thought to be a gas puff as the SOLPS-ITER scenario in question was actually part of a scan for the gas puff in TCV. However, as was indicated in section 4.2, the gas puff only had an effect in the direct vicinity of the divertor target—similar to the effect of parameters like the pitch angle θ, the recycling coefficient and the neutral residence time. Tuning these parameters had some effect on the target plasma and neutral densities but affected little else.

The main other difference that could be observed was that the heat flux in DIV1D was not commensurate with that in SOLPS-ITER. The SOLPS-ITER output showed a rapid decrease of the parallel heat flux in middle region of the divertor leg. From a 2PM point-of-view that was odd, because that postulates that almost all heat is lost directly in front of the target. Although DIV1D is not as strongly simplified as the 2PM, the intuition that most of the power is lost directly in front of the target is still correct. This is related to the fact that the power loss channels in DIV1D—impurity radiation, ionisation and excitation, charge exchange, and recombination—are only relevant in low-temperature regions. In DIV1D, those regions are reserved exclusively for the region directly in front of the target or near the detachment front. As such, no significant power losses are expected outside these regions. The question was what could cause these power losses outside the effects discussed above, and whether including that effect in DIV1D could bring the outputs of the code closer to that of SOLPS-ITER.

It was hoped that by reconstructing the SOLPS-ITER power balance it would be possible to determine those differences. However, this proved outside the scope of this thesis as understanding the workings

of SOLPS-ITER turned out to take more time than was available. While reading up on the SOLPS-ITER code, one thing that kept popping up was the inclusion of experimentally determined anomalous transport coefficients for the perpendicular transport in the code. These were briefly discussed while sketching the background of SOLPS-ITER in section 2.2. While we cannot show definitive proof that it is actually these radial terms that cause the observed losses—attempts at reconstructing the power balance seemed to indicate so but the fact that those reconstructions are not included in this thesis means that they cannot be put forward as evidence—there is the fact that those loss factors are present in SOLPS-ITER (Schneider *et al*, 2006).

Moreover, simulations that were performed after implementing an extra tuneable loss factor in DIV1D—dubbed radial losses and discussed in 4.3—show that the inclusion of an extra loss factor to the power balance allows for DIV1D to yield results that are much closer to the results in SOLPS-ITER than was achievable without such an extra loss factor. This demonstrates, at the very least, that a strong additional loss factor is present in SOLPS-ITER and that it can be reproduced in DIV1D. Preliminary understanding points to perpendicular losses as the culprit, but additional research would need to be performed to establish that for certain.

6.3 ELM-induced dynamics

Chapter 5 dealt with the question of dynamics. Specifically, it looked into the question of what happened in the divertor leg when the plasma was subjected to edge-localised modes on the upstream side. Before it was possible to study these effects, it was necessary to implement an ELM functionality in DIV1D. First, this involved determining the key parameters that should be included in such a description. In following Hill (1997) section 5.1 discussed that those were the total energy content that an ELM dumps into the SOL, i.e. the upstream fluence, the ELM rise and fall times, and the ELM repetition rate. A simple model by Eich *et al* (2017) was introduced that incorporated these key parameters. Truth be told, it did not explicitly mention the repetition rate—as the text is primarily involved with Type I ELMs in which any given ELM will have long since passed before the next one presents itself—but it can be trivially adapted to incorporate that effect.

An alternative formulation was also presented, but was ultimately not used for the simulations discussed in the remainder of the chapter as its advantages fell short of expectations. It was initially developed in the hopes of reducing the strong temporal oscillations that presented themselves in the wake of an ELM, but it turned out that the shape of the ELMs had nothing to do with those oscillations. Despite the fact that it was not useable for the present simulations, however, the developed model still has a nice advantage over the Eich *et al* model. Besides preserving what is good about that model, it additionally removes the three discontinuities in the first derivatives with respect to time. As such is was included as a slightly off-tangent result of the thesis work.

Having implemented this new ELM feature into the DIV1D code, we were free to start probing the dynamics of the detachment front when subjected to ELMs. As a proving grounds for this new feature, a case study was made of the MAST tokamak. This was a logical choice as SD1D, the code by Dudson *et al* (2019) from which DIV1D borrows heavily, was written with that specific device in mind. The 2019 paper discusses SD1D in the context of that tokamak and, consequently, the early simulations in DIV1D were also performed for MAST like parameters. It stood to reason to extend current work rather than start over from scratch. The question that ensued was what ELM parameters were—if not correct—at least remotely plausible in the context of MAST. Sub section 5.2.1 sketched a back-of-the-envelope

calculation to determine those parameters: a big Type I ELM was determined to have a duration of $0.75\ \mathrm{ms}$, and a fluence of $170\ \mathrm{kJ\ m^{-2}}$. Although these values need to be taken with a grain of salt, they provide an idea of the relevant parameter space.

The main results of the ELM scans were presented in sub sections 5.2.2 and 5.2.3. Both weakly ($0.1\ \mathrm{eV} < T_\mathrm{t} < 10\ \mathrm{eV}$) and fully detached plasmas ($T_\mathrm{t} = 0.1\ \mathrm{eV}$) were probed. Two parameters were of primary interest. One was the ELM target fluence, i.e. the amount of extra energy that ends up at the target as a result of the ELM. The other was the position of the detachment front. In the weakly detached plasmas there was a strong correlation between the target fluence and the upstream fluence. It was a not-quite-linear relationship: for lower upstream fluences in the range between 5–$50\ \mathrm{kJ\ m^{-2}}$ the extra fluence on the target actually exceeded the ELM fluence on the upstream side. This suggests that an ELM in this range causes reattachment of the plasma. This is related to a stark increase in the target temperature. This temperature, as we discussed on multiple occasions, is the primary knob which affects the effectiveness of the sinks in the power balance. As the temperature increases, the sinks become less effective and part of the steady-state heat flux which was radiated away in the detached regime reaches the target unmitigatedly after the plasma has reattached.

Oddly, this behaviour does not persist. As the upstream fluence is increased beyond $50\ \mathrm{kJ\ m^{-2}}$ the ratio between target and upstream fluence drops below unity. Section 5.2.2 linked this to the onset of full detachment in the wake of the ELM. The reason for the onset of detachment was shown to be a temporal oscillation in the plasma that persisted long after the ELM had passed. It was very similar to the damped oscillations that was shown in chapter 3 in the context of the artificial viscosity. It is not clear what exactly governs this behaviour, which is an issue as it precludes a clear answer to the question why the plasma detaches in the wake of an ELM for high upstream fluences. Another unexplained effect in this context is the question why—once detachment occurs—the detachment front is suddenly accelerated towards the upstream as can be seen in figure 5.6.

Moving into the results for the fully detached plasma, we additionally observed that there was a threshold value for the ELM fluence below which single-pulse ELMs could not cause reattachment (figure 5.7). For a plasma with $n_\mathrm{u} = 5.0 \times 10^{19}\ \mathrm{m^{-3}}$, $L = 20\ \mathrm{m}$ and $q_\mathrm{ss} = 50\ \mathrm{MW\ m^{-2}}$ a 20-kJ-m^{-2} ELM would not cause reattachment while a 50-kJ-m^{-2} resulted in a fluence at the target of $20\ \mathrm{kJ\ m^{-2}}$and a 140-kJ-m^{-2} ELM resulted in $85\ \mathrm{kJ\ m^{-2}}$ at the target. It was then shown for a 140-kJ-m^{-2} single-pulse ELM, which was previously shown to effect reattachment, that the amount of ELM fluence dropped linearly with increasing upstream density (figure 5.8). This was linked to the fact that the detachment front also linearly receded from the target as a function of increasing upstream density (figure 5.9). Finally, it was shown that for an ELM that could not effect attachment single-pulse, c.q. the 20-kJ-m^{-2} ELM, could cause reattachment in a series of pulses. In the range 250-$1000\ \mathrm{Hz}$, which is outside the range in which individual ELMs overlap on the upstream side but within the repetition rate described by Hill (1997) for Type II and Type III ELMs, a frequency threshold was found between $400\ \mathrm{Hz}$ and $500\ \mathrm{Hz}$. At $400\ \mathrm{Hz}$ the fraction of the upstream fluence that reaches the target is negligible, while it increases to 20% at $500\ \mathrm{Hz}$. Between $500\ \mathrm{Hz}$ and $1000\ \mathrm{Hz}$ the fraction is seen to linearly increase.

In the end, it is not so much these numbers themselves that matter, however, as the fact that they can be obtained using DIV1D. Numbers such as these could be validated in fusion devices[13] to determine whether the DIV1D model holds up to experimental reality.

[13] This is highly unlikely to be in MAST since that device ceased operation in 2013. Quite recently, in the autumn of 2020, its successor MAST Upgrade commenced operation (EUROfusion, 2020).

7. Conclusions and outlook

This thesis sought to answer two research questions related to studying divertor detachment in 1D. Firstly, with respect to the question:

> What physics mechanisms drive the differences between DIV1D and SOLPS-ITER, and can they be accounted for in 1D?

the answer in this thesis is that it is possible—by including an effective radial loss term to the power balance equations of DIV1D—to reproduce the strong power losses that were observed in a SOLPS-ITER simulation of TCV but which could not initially be reproduced in DIV1D. Different loss profiles were tested in the code and their effects on the DIV1D solutions documented. The results showed that the inclusion of these loss profiles allowed for DIV1D solutions to be brought in the range of those obtained in SOLPS-ITER. Crucially, it allowed for a comparable qualitatively similar upstream and target temperature at a comparable upstream density, connection length and incoming heat flux—something that was demonstrated to be impossible without radial losses. However, further research is required to demonstrate beyond doubt that radial losses are actually the cause of this power sink as this thesis failed to provide conclusive evidence for that. Secondly, with respect to the question:

> What is the effect of ELMs on the dynamical behaviour of radiation fronts in the scrape-off layer?

it was found that—using a newly implemented ELM feature—DIV1D is able to simulate the dynamical behaviour of the plasma during and after the passing of an ELM or a series of ELMs. It was established that for a given divertor plasma it is possible to use DIV1D to calculate target fluences and establish whether or not a plasma reattaches. In case of weakly detached plasmas with a steady-state target temperature between $0.1\ \mathrm{eV}$ and $10\ \mathrm{eV}$ ELMs with a given upstream fluence cause an almost *ditto* target fluence. Beyond a certain upstream fluence (dependent on the divertor leg parameters) detachment is observed to ensue in the wake of an ELM. It is unclear at this moment whether this constitutes physically realistic behaviour. Fully detached plasmas with a steady-state target temperature of $0.1\ \mathrm{eV}$ can act as buffers to ELMs up to a certain point. Thresholds, either with respect to the upstream fluence or the repetition frequency, were observed that ELMs need to exceed to cause reattachment. Finally, for a given upstream fluence the target fluence was found to be linearly dependent on the upstream density.

The research questions now having been answered to the extent allowed by the present work, this really only leaves one last question to be tackled. What would be interesting to ask of future research? The rest of this chapter is dedicated to some thoughts on this.

> A better grasp of what happens under the hood of DIV1D at the exact moment that detachment occurs would be well worth the effort.

The DIV1D code works, and it works nicely at that. However, throughout this work there were several occasions on which we had to ask ourselves whether what we were seeing was physical. This started with a strange 'kink' in the solutions just in front of the divertor target in chapter 3. Although it was discussed that the effect of such a kink was most likely limited, it still looked off. In chapter 5 it was observed that despite the plasma being fully detached, erratic temperature spikes could sometimes be

observed at the divertor target during periods in which the temperature should be minimal. Again, this seemed to have little effect on the solutions but it looked strange. Furthermore, when detachment occurred the detachment front was suddenly seen to accelerate towards the upstream side of the leg. Why? Additionally, the code got extremely slow when doing calculations for detached plasma as compared to attached cases. Future work to understand what happens in DIV1D at the exact moment when full detachment occurs would be a great way to get more confidence in the code. This is particularly important when DIV1D is used to simulate the dynamics of the divertor leg subjected to ELMs, as this project showed that detachment and reattachment can happen in these cases.

> The nature of the decaying oscillation that features almost universally in DIV1D solutions should be established: is it a numerical artefact or a physically interpretable factor?

A major driver for detachment in the wake of ELMs was the strong oscillatory behaviour in many of the plasma parameters. Since this behaviour was not limited to ELMs, though, also appearing during the artificial viscosity and radial losses scans, it deserves its own recommendation. Some effort was spent on determining the nature of this oscillation, but no clear conclusions could be drawn. As this feature is very typical of DIV1D simulations, it seems prudent to try and find out what causes it. This is particularly pressing due to the previously discussed upstream acceleration of the detachment front after detachment: accessing the fully detached 0.1-eV regime on the target side has been shown to have large consequences for the state of the plasma, and hence any factor that contributes significantly to accessing this regime should be well understood.

> To determine whether the ELM-induced dynamics of DIV1D have any bearing on reality, it would be interesting to compare them to experiments.

The behaviour of the divertor plasma when subjected to ELMs has been probed in this thesis via a case study of MAST. There is really no way to determine whether the answers currently provided by DIV1D hold up to reality. It would thus be informative to compare DIV1D to experimental results to get an understanding of what, if any, aspects of the dynamics are correctly predicted by DIV1D and what aspects require further work.

> A dedicated benchmark of DIV1D against 2D SOLPS-ITER simulations should be performed.

It would be highly commendable to perform such a benchmark, because if DIV1D is to be of use to simulate the dynamics of the divertor leg it should be able to at least qualitatively reproduce the results of more all-encompassing models. A limited attempt at a benchmark was made during this project, and if that showed one thing it was that SOLPS-ITER is very complex. It is advisable that any such dedicated benchmark would be performed under the supervision of someone with a strong basis in SOLPS-ITER. One interesting suggestion by J. Koenders was to benchmark DIV1D not against a single flux tube but rather against a 'averaged flux tube', i.e. which is radially averaged over several flux tubes.

> The role of radial losses in the onset of detachment might be a non-negligible factor, which makes it a potential candidate for further study.

Chapter 4 indicates that radial power losses might play a role in the onset of detachment. It would be interesting to subject that hypothesis to further study.

Bibliography

Bethe, H.A. & Bacher, R.F. (1936), 'A. Stationary States of Nuclei', *Reviews of Modern Physics*, 8(2) [online]. Available at: https://doi.org/10.1103/RevModPhys.8.82 (Accessed: 12 November 2020)

Bethe, H.A. (1937) 'B. Nuclear Dynamics, Theoretical', *Reviews of Modern Physics*, 9(2) [online]. Available at: https://doi.org/10.1103/RevModPhys.9.69 (Accessed: 12 November 2020)

Bonnin, X. & Pitts, R. (2015) *ITER Newsline: ITER unveils a new tool for plasma edge modelling: SOLPS-ITER* [online]. Available at: https://www.iter.org/newsline/-/2168 (Accessed: 5 October 2020)

Bonnin, X., Pitts, R.A., Komarov, V., Escourbiac, F., Merola, M., Bo, L., Wei, L., Pan, L. & Kukushkin, A.S. (2017) 'ITER divertor plasma response to time-dependent impurity injection', *Nuclear Materials and Energy*, 12 [online]. Available at: https://doi.org/10.1016/j.nme.2017.03.010 (Accessed: 16 October 2020)

Boozer, A.H. (2018) 'Enhanced control', *Nature Physics,* 14(12) [online]. Available at: https://doi.org/10.1038/s41567-018-0281-y (Accessed: 2 October 2020)

Braginskii, S.I. (1965) 'Transport processes in a plasma', *Reviews of plasma physics*, 1 [online]. Available at: https://static.ias.edu/pitp/2016/sites/pitp/files/braginskii_1965-1.pdf (Accessed: 2 October 2020)

Cohen, R.S., Spitzer, Jr., L. & Routly, P.McR. (1950) 'The Electrical Conductivity of an Ionized Gas', *Physical Review,* 80(2) [online]. Available at: https://doi.org/10.1103/PhysRev.80.230 (Accessed: 20 January 2021)

Dudson, B.D., Allen, J., Body, T., Chapman, B., Lau, C., Townley, L., Moulton, D., Harrison, J. & Lipschultz, B. (2019) 'The role of particle, energy and momentum losses in 1D simulations of divertor detachment', *Plasma Physics and Controlled Fusion*, 61(6) [online]. Available at: https://doi.org/10.1088/1361-6587/ab1321 (Accessed: 4 December 2019)

Dudson, B. (2016) *SD1D: 1D divertor model for detachment studies* [online]. Available at: https://github.com/boutproject/SD1D/blob/master/doc/sd1d-manual.tex (Accessed: 5 January 2021)

Eddington, A.S. (1920) 'The Internal Constitution of the Stars', *The Scientific Monthly,* 11(4) [online]. Available at: http://www.jstor.com/stable/6491 (Accessed: 7 September 2020)

Eich, T., Sieglin, B., Thornton, A.J., Faitsch, M., Kirk, A., Herrmann, A., Suttrop, W., JET contributors, MST contributors, ASDEX Upgrade & MAST teams (2017) 'ELM divertor peak energy fluence scaling to ITER with data from JET, MAST and ASDEX upgrade', *Nuclear Materials and Energy*, 12 [online]. Available at: https://doi.org/10.1016/j.nme.2017.04.014 (Accessed: 20 March 2020)

Einstein, A. (1905) 'Ist die Trägheit eines Körpers von seinem Energieinhalt abhängig?', *Annalen der Physik*, 323(13) [online]. Available at: https://doi.org/10.1002/andp.19053231314 (Accessed: 12 November 2020)

EUROfusion (2018) *European Research Roadmap to the Realisation of Fusion Energy*. Garching: EUROfusion [online]. Available at: https://www.euro-fusion.org/eurofusion/roadmap/ (Accessed: 29 September 2020)

EUROfusion (2020) *MAST Upgrade starts up its first plasma*. Available at: https://www.euro-fusion.org/news/2020/october/mast-upgrade-starts-up-its-first-plasma/ (Accessed: 22 January 2021)

Field, A.R., Balboa, I., Drewelow, P., Flanagan, J., Guillemaut, C., Harrison, J.R., Huber, A., Huber, V., Lipschultz, B., Matthews, G., Meigs, A., Schmitz, J., Stamp, M., Walkden, N. & JET contributors (2017) 'Dynamics and stability of divertor detachment in H-mode plasmas on JET', *Plasma Physics and Controlled Fusion*, 59(9) [online]. Available at: https://doi.org/10.1088/1361-6587/aa764c (Accessed: 16 October 2020)

Fourier, J.B.J. (1822) *Théorie analytique de la chaleur*. Paris: Firmin Didot [online]. Available at: https://fr.wikisource.org/wiki/Th%C3%A9orie_analytique_de_la_chaleur# (Accessed: 22 December 2020)

Freidberg, J. (2007) *Plasma Physics and Fusion Energy*. Cambridge: Cambridge University Press.

Greenwald, M. (2002) 'Density limits in toroidal plasmas', *Plasma Physics and Controlled Fusion,* 44(8) [online]. Available at: https://doi.org/10.1088/0741-3335/44/8/201 (Accessed: 29 September 2020)

Harrison, J.R., Vijvers, W.A.J., Theiler, C., Duval, B.P., Elmore, S., Labit, B., Lipschultz, B., Van Limpt, S.H.M., Lisgo, S.W., Tsui, C.K., Reimerdes, H., Sheikh, U., Verhaegh, K.H.A., Wischmeier, M. & the MST1 and TCV Teams (2017) 'Detachment evolution on the TCV tokamak', *Nuclear Materials and Energy*, 12 [online]. Available at: https://doi.org/10.1016/j.nme.2016.10.020 (Accessed: 28 September 2020)

Helander, P., Beilder, C.D., Bird, T.M., Drevlak, M., Feng, Y., Hatzky, R., Jenko, F., Kleiber, R., Proll, J.H.E., Turkin, Y. & Xanthopoulos, P. (2012) 'Stellarator and tokamak plasmas: a comparison', *Plasma Physics and Controlled Fusion*, 54(12) [online]. Available at: https://doi.org/10.1088/0741-3335/54/12/124009 (Accessed: 5 January 2021)

Hill, D.N. (1997) 'A review of ELMs in divertor tokamaks', *Journal of Nuclear Materials*, 241–243 [online]. Available at: https://doi.org/10.1016/S0022-3115(97)80039-6 (Accessed: 13 October 2020)

IPCC (2014) 'Summary for Policymakers' in Core Writing Team, Pachauri, R.K. & Meyer, L.A. (eds.) *Climate Change 2014: Synthesis Report. Contribution of Working Groups I, II and III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. Geneva: IPCC [online]. Available at: https://www.ipcc.ch/report/ar5/syr/ (Accessed: 15 October 2020)

IPCC (2018), 'Summary for Policymakers' in Masson-Delmotte, V., Zhai, P., Pörtner, H.-O., Roberts, D., Skea, J., Shukla, P.R., Pirani, A., Moufouma-Okia, W., Péan, C., Pidcock, R., Connors, S., Matthews, J.B.R., Chen, Y., Zhou, X., Gomis, M.I., Lonnoy, E., Maycock, T., Tignor M. & Waterfield, T. (eds.) *Global Warming of 1.5°C. An IPCC Special Report on the impacts of global warming of 1.5°C above pre-industrial levels and related global greenhouse gas emission pathways, in the context of strengthening the global response to the threat of climate change, sustainable development, and efforts to eradicate poverty*. In Press [online]. Available at: https://www.ipcc.ch/sr15/ (Accessed: 15 October 2020)

ITER Organization (2020), *Home page* [online]. Available at: https://www.iter.org/ (Accessed: 12 November 2020)

ITER Physics Expert Group on Confinement and Transport, ITER Physics Expert Group on Confinement Modelling and Database & ITER Physics Basis Editors (1999) 'Chapter 2: Plasma Confinement and transport', *Nuclear Fusion,* 39(12) [online]. Available at: https://doi.org/10.1088/0029-5515/39/12/302 (Accessed: 24 January 2021)

ITER Physics Expert Group on Divertor, ITER Physics Expert Group on Divertor Modelling and Database & ITER Physics Basis Editors (1999) 'Chapter 4: Power and particle control', *Nuclear Fusion*, 39(12) [online]. Available at: https://doi.org/10.1088/0029-5515/39/12/304 (Accessed: 13 December 2020)

Kabinet-Rutte III (2019) *Klimaatakkoord* [online].
Available at: https://www.klimaatakkoord.nl/documenten/publicaties/2019/06/28/klimaatakkoord (Accessed: 12 November 2020)

Koenders, J. (2020) *Towards model-based control of divertor detachment: Control-oriented modelling and dynamical analysis on TCV*, Master's thesis. Eindhoven: Eindhoven University of Technology.

Kirk, A., Counsell, G.F., Wilson, H.R., Ahn, J.-W., Akers, R., Arends, E.R., Dowling, J., Martin, R., Meyer, H. & Hole, M. (2004) 'ELM characteristics in MAST', *Plasma Physics and Controlled Fusion*, 46(3) [online]. Available at: https://doi.org/10.1088/0741-3335/46/3/009 (Accessed: 5 January 2021)

Lawson, J.D. (1957) 'Some Criteria for a Power Producing Thermonuclear Reactor', *Proceedings of the Physical Society. Section B*, 70(1) [online]. Available at: https://doi.org/10.1088/0370-1301/70/1/303 (Accessed: 12 November 2020)

Lieberman, M.A. & Lichtenberg, A.J. (2005) *Principles of Plasma Discharges and Materials Processing. Second Edition.* Hoboken: John Wiley & Sons, Inc.

Livingston, M.S. & Bethe, H.A. (1937), 'C. Nuclear Dynamics, Experimental', *Reviews of Modern Physics*, 9(3) [online]. Available at: https://doi.org/10.1103/RevModPhys.9.245 (Accessed: 12 November 2020).

Lopes Cardozo, N. (2018) *Fusion on the back of an envelope*, lecture notes, Fusion on the back of an Envelope 3MF100, Eindhoven University of Technology, delivered autumn 2018.

Nakazawa, S., Nakajima, N., Okamoto, M. & Ohyabu, N. (2000) 'One-dimensional simulation on stability of detached plasma in a tokamak divertor', *Plasma Physics and Controlled Fusion*, 42(4) [online]. Available at: https://doi.org/10.1088/0741-3335/42/4/303 (Accessed: 9 January 2020)

Maisonnier, D., Cook, I., Sardain, P., Boccaccini, L, Bogusch, E., Broden, K., Di Pace, L., Forrest, R., Giancarli, L., Hermsmeyer, S., Nardi, C., Norajitra, P., Pizzuto, A., Taylor, N. & Ward, D. (2005) 'The European power plant conceptual study', *Fusion Engineering and Design,* 75–79 [online]. Available at: https://doi.org/10.1016/j.fusengdes.2005.06.095 (Accessed: 7 October 2020)

MDSPlus (2020) *Home page* [online]. Available at: https://www.mdsplus.org/index.php/Introduction (Accessed: 16 December 2020)

Meyer, H. (2016) 'The mega amp spherical tokamak', in Neilson, G.H. (ed.) *Magnetic Fusion Energy*. Cambridge: Woodhead Publishing [online]. Available at: https://doi.org/10.1016/B978-0-08-100315-2.00012-X (Accessed: 26 October 2020)

Michoski, C., Dawson, C., Kubatko, E.J., Wirasaet, D., Brus, S. & Westerink, J.J. (2016) 'A Comparison of Artificial Viscosity, Limiters, and Filters, for High Order Discontinuous Galerkin Solutions in Nonlinear Settings', *Journal of Scientific Computing*, 66 [online]. Available at: https://doi.org/10.1007/s10915-015-0027-2 (Accessed: 22 January 2021)

Pitts, R.A., Bardin, S., Bazulev, N., Van den Berg, M.A., Bunting, P., Carpentier-Chouchana, S., Coenen J.W., Corre, Y., Dejarnac, R., Escourbiac, F., Gaspar, J., Gunn, J.P., Hirai, T., Hong S-H., Horacek, J., Iglesias, D., Komm, M., Krieger, K., Lasnier, C., Matthews, G.F., Morgan, T.W., Panayotis, S., Pestchanyi, S., Podolnik, A., Nygren, R.E., Rudakov, D.L., De Temmerman, G., Vondracek, R. & Watkins, J.G. (2017) 'Physics conclusions in support of ITER W divertor monoblock shaping', *Nuclear Materials and Energy,* 12 [online]. Available at: https://doi.org/10.1016/j.nme.2017.03.005 (Accessed: 7 October 2020).

Pitts, R.A., Bonnin, X., Escourbiac, F., Frerichs, H., Gunn, J.P., Hirai, T., Kukushkin, A.S., Kaveeva, E., Miller, M.A., Moulton, D., Rozhansky, V., Senichenkov, I. Sytova, E., Schmitz, O., Stangeby, P.C., De Temmerman, G., Veselova, I. & Wiesen, S. (2019) 'Physics basis for the first ITER tungsten divertor', *Nuclear Materials and Energy*, 20 [online]. Available at: https://doi.org/10.1016/j.nme.2019.100696 (Accessed: 7 October 2020)

Post, D.E., Jensen, R.V., Tarter, C.B., Grasberger, W.H. & Lokke, W.A. (1977) 'Steady-state radiative cooling rates for low-density, high-temperature plasmas', *Atomic Data and Nuclear Data Tables*, 20(5) [online]. Available at: https://doi.org/10.1016/0092-640X(77)90026-2 (Accessed: 7 January 2020)

Reiter, D. (2019) *The EIRENE Code User Manual*. Available at: http://eirene.de/html/manual.html (Accessed: 5 January 2021)

Rocco, P. & Zucchetti, M. (2000) 'Waste management for different fusion reactor designs', *Journal of Nuclear Materials*, 283–287(2) [online]. Available at: https://doi.org/10.1016/S0022-3115(00)00127-6 (Accessed: 12 November 2020)

Rognlien, T.D., Umansky, M.V., Xu, X.Q. & Cohen, R.H. (2004) 'Self-consistent simulation of turbulence and transport in tokamak edge plasmas', *Contributions to Plasma Physics*, 44(1–3) [online]. Available at: https://doi.org/10.1002/ctpp.200410026 (Accessed: 5 October 2020)

Saarinen, L., Dahlbäck, N. & Lundin, U. (2015) 'Power system flexibility need induced by wind and solar power intermittency on time scales of 1–14 days', *Renewable Energy*, 83 [online]. Available at: https://doi.org/10.1016/j.renene.2015.04.048 (Accessed: 12 November 2020)

Schneider, R., Bonnin, X., Borrass, K., Coster, D.P., Kastelewicz, H., Reiter, D., Rozhansky, V.A. & Braams, B.J. (2006) 'Plasma Edge Physics with B2-Eirene', *Contributions to Plasma Physics*, 46(1–2) [online]. Available at: https://doi.org/10.1002/ctpp.200610001 (Accessed: 30 June 2020)

Shafranov, V.D. (2001) 'The initial period in the history of nuclear fusion research at the Kurchatov Institute', *Physics – Uspekhi*, 44(8) [online].
Available at: https://doi.org/10.1070/PU2001v044n08ABEH001068 (Accessed: 6 October 2020)

Spitzer, Jr., L. & Härm, R. (1953) 'Transport Phenomena in a Completely Ionized Gas', *Physical Review*, 89(5) [online]. Available at: https://doi.org/10.1103/PhysRev.89.977 (Accessed: 20 January 2021)

Stangeby, P.C. (2000) *The Plasma Boundary of Magnetic Fusion Devices.* New York: Taylor & Francis Group [online]. Available at: https://doi.org/10.1201/9780367801489 (Accessed: 6 December 2019)

Stangeby, P.C. (2018) 'Basic physical processes and reduced models for plasma detachment', *Plasma Physics and Controlled Fusion*, 60(4) [online]. Available at: https://doi.org/10.1088/1361-6587/aaacf6 (Accessed: 9 December 2019)

Thacker, B.H., Doebling, S.W., Hemez, F.M., Anderson, M.C., Pepin, J.E. & Rodriguez, E.A. (2004) *Concepts of Model Verification and Validation*. Los Alamos: Los Alamos National Laboratory [online]. Available at: https://doi.org/10.2172/835920 (Accessed: 22 January 2021)

Thornton, A.J., Kirk, A. & The MAST Team (2014) 'Scaling of the scrape-off layer width during inter-ELM H modes on MAST as measured by infrared thermography', *Plasma Physics and Controlled Fusion*, 56(5) [online]. Available at: https://doi.org/10.1088/0741-3335/56/5/055008 (Accessed: 26 October 2020)

Van Zalk, J. & Behrens, P. (2018) 'The spatial extent of renewable and non-renewable power generation: A review and meta-analysis of power densities and their application in the U.S.', *Energy Policy*, 123 [online]. Available at: https://doi.org/10.1016/j.enpol.2018.08.023 (Accessed: 12 November 2020)

Veres, G., Pitts, R.A., Bencze, A., Márki, J., Tál, B., Tye, R & TCV Team (2009) 'Fast radiation dynamics during ELMs on TCV', *Journal of Nuclear Materials*, 390–391 [online]. Available at: https://doi.org/10.1016/j.jnucmat.2009.01.220 (Accessed: 19 March 2020)

Von Neumann, J. & Richtmyer, R.D., 'A Method for the Numerical calculation of Hydrodynamic Shocks', *Journal of Applied Physics*, 21(3) [online]. Available at: https://doi.org/10.1063/1.1699639 (Accessed: 22 January 2021)

Wagner, F., Becker, G., Behringer, K., Campbell, D., Eberhagen, A., Engelhardt, W., Fussmann, G., Gehre, O., Gernhardt, J., Gierke, G. v., Haas, G., Huang, M., Karger, F., Keilhacker, M., Klüber, O., Kornherr, M., Lackner, K., Lisitano, G., Lister, G.G., Mayer, H.M., Meisel, D., Müller, E.R., Murmann, H., Niedermeyer, H., Poschenrieder, W., Rapp, H., Röhr, H., Schneider, F., Siller, G., Speth, E., Stäbler, A., Steuer, K.H., Venus, G., Vollmer, O. & Yü, Z. (1982) 'Regime of Improved Confinement and High Beta in Neutral-Beam-Heated Divertor Discharges of the ASDEX Tokamak', *Physical Review Letters*, 49(19) [online]. Available at: https://doi.org/10.1103/PhysRevLett.49.1408 (Accessed: 19 September 2019)

Wensing, M. (2020). Email to Jens Peter Frankemölle, February 18.

Wensing, M., Loizu, J., Reimerdes, H., Duval, B.P., Wischmeier, M. and the TCV team (2020) 'X-point potential well formation in diverted tokamaks with unfavorable magnetic field direction', *Nuclear Fusion,* 60(5) [online]. Available at: https://doi.org/10.1088/1741-4326/ab7d4f (Accessed: 19 January 2021)

Westerhof, E. (2019) *DIV1D manual: a 1D code for divertor plasma simulation.* Eindhoven: DIFFER [unpublished].

Zhang, D.R., Chen, Y.P., Xu, X.Q. and Xia, T.Y. (2019) 'Self-consistent simulation of transport and turbulence in tokamak edge plasma by coupling SOLPS-ITER and BOUT++', *Physics of Plasmas*, 26(1) [online]. Available at: https://doi.org/10.1063/1.5084093 (Accessed: 2 October 2020)

Zhang, D.R., Chen, Y.P., Xu, X.Q. and Xia, T.Y., Liu, S.C. & EAST team (2020) 'Simulation of EAST edge plasma using SOLPS-ITER/BOUT++ coupling', *Nuclear Fusion*, 60(10) [online]. Available at: https://doi.org/10.1088/1741-4326/abaa90 (Accessed: 5 October 2020)

Photographic material

Figure 1.1 is an adaptation from: Rswilcox (2017) 'A worker inside the DIII-D vacuum vessel during a maintenance period in 2017', *Wikimedia Commons*.
Available at: https://commons.wikimedia.org/wiki/File:2017_TOCAMAC_Fusion_Chamber_N0689.jpg under a CC BY-SA 4.0 license

Appendices

A. Fortran implementation of radial losses in DIV1D

Introducing losses perpendicularly to the flux tube was a rather straightforward exercise. It involved adding a Fortran subroutine, dubbed calculate_radial_losses, to the existing source code. This subroutine is called as part of the calculate_sources subroutine—which, as its name suggests, calculates the source and sink terms discussed in section 2.3—and is in turn called by the right_hand_side subroutine. Finally, the right_hand_side subroutine is called in the main program at each internal time step of the solver to supply it with the equations to be solved.

Box A.1. Calculate_radial_losses subroutine

```
subroutine calculate_radial_losses(Nx,radial_sink,q_parallel)

        ! This subroutine captures the radial losses as a volumetric energy sink with a gaussian
        ! profile. Inputs are the gaussian width and peak location, given by radial_loss_width and
        ! radial_loss_location respectively. Depending on whether radial_loss_gaussian is zero or
        ! not, the radial loss profile is a bell curve or constant, respectively. The normalisation
        ! is calculated numerically, so that the total lost heat flux is always fixed by q_parX.

        implicit none
        integer :: Nx
        real(wp) :: radial_sink(Nx), a0, x0, norm, gaussian(Nx), normalisation,q_parallel(Nx)

        if (radial_loss_gaussian.gt.0) then
                a0 = radial_loss_width
                x0 = radial_loss_location
                gaussian = exp(-(x-x0)**2/(2*a0**2))
                normalisation = sum(gaussian * delta_xcb)
                radial_sink = radial_loss_factor * q_parX * gaussian / normalisation
        elseif (radial_loss_gaussian.lt.0) then
                radial_sink = radial_loss_factor *q_parallel / L
        else
                radial_sink = radial_loss_factor * q_parX / L
        endif

end subroutine calculate_radial_losses
```

Box A.1 shows the calculate_radial_losses subroutine as it has been implemented in the DIV1D code. It requires an additional four global parameters on the input side. These have been written down in box A.3. Depending on the shape of the loss profile—defined by radial_loss_gaussian—the radial_loss_factor takes on an absolute or relative meaning. If the user sets radial_loss_gaussian to 0 (constant losses) or any positive number (gaussian losses), radial_loss_factor directly reflects the fraction of the upstream heat flux (q_parX) that is lost over the length of the flux tube. If instead radial_loss_gaussian is set to a negative value, radial_loss_factor reflects the fraction of the local heat flux density q_parallel(ix)/L that is lost in a given grid cell. In total, the amount of lost heat flux is lower since for most of the flux tube it holds that q_parallel(ix) $<$ q_parX. Finally, in case a gaussian radial loss profile is selected then the user is also required to supply the position of peak losses radial_loss_location and the width of the gaussian radial_loss_width.

Irrespective of the type of loss profile that is selected, the sink terms are calculated at each position in the flux tube and are saved to an array called radial_sink, which is outputted by the subroutine. The calculate_radial_losses subroutine is consecutively called in the pre-existing calculate_sources subroutine. As inputs, outside the global ones provided in box A.3 plus an array with the width of each grid cell delta_xcb (required for numerically normalising the gaussian), it additionally requires the grid size Nx and the current iteration of q_parallel.

Finally, an earlier iteration of the code would crash whenever the plasma went into full detachment. After some troubleshooting, it was realised that the culprit was that in cases of full detachment the heat flux would dip below zero: while all the other sinks went to zero as detachment happened the radial losses did not. To ameliorate this issue, after calling the calculate_radial_losses subroutine, a check was included to make sure that whenever the radial sink becomes bigger than 90% of the residual heat flux over the cell width, q_parallel/delta_xcb, the radial sink is set to zero for every position from there until the target. Having added this additional check, the code also converged in fully detached cases.

Box A.2. Addition to calculate_sources subroutine

```
! Some additional parameters needed to be declared at the beginning of the subroutine

real(wp) :: radial_sink(Nx)
integer  :: ix, iix

! The first part of this subroutine, omitted here, calls all other loss terms and lumps them in the array Source_Q(Nx).

call calculate_radial_losses(Nx,radial_sink,q_parallel)
! Consecutively, check whether subtracting this radial_sink does not yield
! unphysical results by confirming that the total losses over the flux tube are
! smaller than the incoming flux, so as not to obtain sub-zero fluxes.
do ix = 1, Nx
        if (abs(Source_Q(ix) - radial_sink(ix)).ge.0.9d+0*abs(q_parallel(ix)/delta_xcb(ix))) then
                do iix = ix, Nx
                        radial_sink(iix) = 0
                enddo
                exit
        endif
enddo
Source_Q = Source_Q - radial_sink

! The last bit of the subroutine, likewise omitted, offers an option to call a spike filter.
```

Box A.3. Global parameters in physics_parameters module

```
integer, parameter, private :: wp = KIND(1.0D0)

real(wp)  :: radial_loss_factor     = 0      ! percentage of the parallel flux that is lost radially throughout the flux tube
Integer   :: radial_loss_gaussian   = 0    ! set to 0 (default) for a constant loss factor, or to 1 for a gaussian distribution
real(wp)  :: radial_loss_width      = 1d+20
! determine width of radial loss distribution (only used for radial_loss_gaussian = 1)
real(wp)  :: radial_loss_location   = 0
! determine peak location of radial loss distribution (only used for radial_loss_gaussian = 1)
```

Appendix B: Fortran implementation of ELMs in DIV1D

To incorporate the effect of ELMs in DIV1D a set of three subroutines was written. The user interacts with these subroutines through a set of input parameters that are defined globally. These parameters have been included in box B.1. The consist of a set of three switches with which the user decides what type of ELM is wanted. Depending on what switches are turned on the user is required to set additional parameters. Regardless of which switch is turned on, the user is always advised to explicitly set elm_start_time, which determines at what outer time step the (first) ELM starts, and elm_ramp_time, which determines how fast the ELM reaches its peak value. It is related to the variable τ_r in equations 5.1 and 5.5–5.6. Note: all time-related parameters are defined in units of time step size so to obtain the time in seconds one needs to multiply those values by delta_t, the outer time step, which is also an input parameter. Likewise, the user is always advised to explicitly set gaussian_elm to either one or not one, which triggers the subroutine to employ gaussian or triangular ELMs respectively. More on that in a second.

Whether the other parameters need to be defined depends on what switch(es) is or are activated. If switch_elm_heat_flux is set to one, the user also has to input elm_expelled_heat, which is a slight misnomer that refers to the ELM fluence, i.e. the parameter Q_{ELM} in chapter 5. If switch_elm_density is set to one, the user also has to input elm_expelled_particles which we will discuss further below. Finally, if switch_elm_series is set to one the user also has to input elm_time_between. This value is related to the ELM repetition frequency f_{ELM} via

$$f_{\mathrm{ELM}}\ [\mathrm{Hz}] = 1/(\mathrm{elm\_time\_between} * \mathrm{delta\_t}). \tag{A.1}$$

All global parameters related to ELMs are included in the physics_parameters module, a part of which has been included as box B.1.

Box B.1. Global ELM parameters in physics_parameters module

```
integer, parameter, private :: wp = KIND(1.0D0)
integer  :: elm_start_time          = 0         ! time step (outer step) at which the ELM starts
integer  :: elm_ramp_time           = 0         ! time (outer step) over which the ELM ramps up
integer  :: elm_time_between        = 200000000 ! time (outer step) between two ELMs
real(wp) :: elm_expelled_heat       = 0.0d+0  ! total heat flux owing to elm, integrated over time [W/m^2]
real(wp) :: elm_expelled_particles  = 0.0d+0  ! total particle flux owing to elm, integrated over time [/m^2]
integer  :: switch_elm_heat_flux    = 0         ! turns off (0, default) or on (1) the elm contribution to the heat flux
integer  :: switch_elm_density      = 0         ! turns off (0, default) or on (1) the elm contribution to the particle flux
integer  :: switch_elm_series       = 0         ! turns off (0, default) or on (1) the multi-elm sequence
integer  :: gaussian_elm            = 1! switch between gaussian ELM (1, default) and triangular ELM (anything but 1)
```

Depending on whether the user selects a gaussian ELM or a triangular ELM, one of two subroutines is called by the main ELM subroutine (simulate_elm in box B.4, which will be discussed in a second). These have been included as boxes B.2 and B.3, respectively. Both subroutines are straightforward implementations of equation 5.1 (triangular model) and equations 5.5–5.6 (gaussian model). The only difference is that, besides $q_{\mathrm{ELM}}(t)$ itself, these subroutines also calculate the first derivative with respect to time, $\mathrm{d}q_{\mathrm{ELM}}/\mathrm{d}t$. This extra step is linked to the reason why the key input parameter to these functions, the 'fluence' has been rather generally dubbed time_integrated_quantity while the output parameters are called quantity and dquantity, referring to $q_{\mathrm{ELM}}(t)$ and its first derivative respectively.

Moving to simulate_elm itself, we can see why these more general names were chosen. While equations 5.1 and 5.5–5.6 were defined with respect to the heat flux, they might just as well be used for the density. ELMs, after all, do not just perturb the power balance. There is also a significant outflux of particles. In DIV1D it was attempted to model this effect by ramping the upstream density over the course of the ELM. DIV1D already incorporated the possibility of ramping the upstream density over the course of time by including a density ramp rate. The word 'rate' here suggests why it was necessary to calculate the first derivative as well: it was not the extra density but the rate of change of the density that DIV1D needed.

Box B.2. Maxwell_boltzmann_elm subroutine

```
subroutine maxwell_boltzmann_elm(time_relative,time_integrated_quantity,quantity,dquantity)

        ! This subroutine calculates the temporal evolution of an ELM at a fixed point in space,
        ! typically at the X-point boundary, as the sum of four Maxwell-Boltzmann distributions,
        ! MB0(t) through MB3(t), which are normalised and then multiplied by the input parameter
        ! time_integrated_quantity. It takes and elm_ramp_time and delta_t to calculate the steepness
        ! of the initial growth. The mb2mb_distance is a fixed scaling factor that helps with keeping
        ! the ELM relatively smooth, minimising bumps in the tail of the profile.
        !
        ! The MBs are used in calculating, for example, the heat flux boundary q_parX in the
        ! right_hand_side subroutine when the time_integrated_quantity is chosen to reflect to total
        ! energy that is dumped into the flux tube. The analytical derivates of the MBs, c.q. the dMBs,
        ! are used in that same subroutine to calculate the density boundary, which requires the
        ! change in density rather than the absolute density, when time_integrated_quantity is set to
        ! the total number of particles dumped into the flux tube.

        implicit none

        real(wp)      :: t, tr, norm, mb2mb_distance
        real(wp)      :: a0, a1, a2, a3, MB0, MB1, MB2, MB3, dMB0, dMB1, dMB2, dMB3
        real(wp)      :: time_relative, time_integrated_quantity, quantity, dquantity
        integer       :: istep

        t = time_relative
        tr = elm_ramp_time*delta_t
        norm   = 1d0/2d0*sqrt(2d0/pi)
        mb2mb_distance = 1.4d0

        a0     = mb2mb_distance**0d0*tr/sqrt(2d0)
        a1     = mb2mb_distance**1d0*tr/sqrt(2d0)
        MB0    = norm/a0**3d0*t**2d0*exp(-t**2d0/(2d0*a0**2d0))
        MB1    = norm/a1**3d0*t**2d0*exp(-t**2d0/(2d0*a1**2d0))
        quantity   = (MB0 + MB1)*time_integrated_quantity

        dMB0   = (2d0/t-t/a0**2d0)*MB0
        dMB1   = (2d0/t-t/a1**2d0)*MB1
        dquantity  = (dMB0 + dMB1)*time_integrated_quantity

end subroutine maxwell_boltzmann_elm
```

The simulate_elm subroutine is also responsible for temporal accounting. Both ELM shape subroutines are just simple functions of time. It is in simulate_elm that the delay due to elm_start_time is accounted for. Likewise, it keeps track of how many ELMs are present in the system at any given time—depending

on the value of elm_time_between, which gives the number of time steps between the starts of two ELMs, multiple ELMs can be present at the same time. The simulate_elm subroutine remembers how many ELMs are present and when they started. This subroutine is called in the right_hand_side (which is called at each iteration of the solver) and the values quantity and dquantity are used to manipulate the boundary conditions at the X-point, which are set in that subroutine. See box B.5.

Box B.3. Triangular_elm subroutine

```
subroutine triangular_elm(time_relative,time_integrated_quantity,quantity,dquantity)

        ! Time-dependent ELM model from Eich et al (2017) [https://doi.org/10.1016/j.nme.2017.04.014]
        ! that simulates ELMs as triangular waveforms.

        implicit none

        real(wp)      :: t, tr, norm, quantity, dquantity, time_relative, time_integrated_quantity

        t = time_relative
        tr = elm_ramp_time*delta_t
        norm = time_integrated_quantity*(2d0/3d0)*(1/tr)

        if (t.lt.tr) then
                quantity = norm * t/tr
                dquantity= norm * 1/tr
        elseif (t.gt.tr.and.t.le.3d0*tr) then
                quantity = norm * (1d0-(t-tr)/(2d0*tr))
                dquantity= norm * -1d0/(2d0*tr)
        endif

end subroutine triangular_elm
```

Box B.4. Simulate_elm subroutine

```
subroutine simulate_elm(elm_heat_load, elm_density_change, time)

! This subroutine translates the ELM that comes out of maxwell_boltzmann_elm() or triangular_elm() into
! parameters that can be added to the boundary conditions for density and heat flux at the X-point
! in the right_hand_side routine. Density and heat fluxes can be turned on independently through a
! pair of switches, which should be set to either 0 or 1. Two if statements test against this prior to execution
! to save computational time. The elm_count and time_mod combined with the while loop act as a counter to keep
! track of how many ELMs are currently in the system. Setting switch_elm_series (accepts 0 and 1) to 0 zero
! effectively kills the while loop at the end of the first iteration, resulting in only a single ELM in that case.

implicit none
! Input parameters
real(wp)              :: time
! Internal parameters
integer               :: elm_count
real(wp)              :: time_relative, time_mod
! Parameters required by maxwell_boltzmann_elm()
real(wp)              :: time_integrated_quantity, quantity, dquantity
! Output parameters
real(wp)              :: elm_heat_load, elm_density_change
```

Box B.4. Simulate_elm subroutine *continued*

```
! Check against illegal input (which only happens if a user actively alters either of the switches,
! since they are set to zero by default in physics_parameters).

if(switch_elm_density.ne.0.and.switch_elm_density.ne.1) then
        print *,"Error: switch_elm_density only accepts 0 or 1 as inputs. Terminating div1d.exe."
        call exit
elseif(switch_elm_heat_flux.ne.0.and.switch_elm_heat_flux.ne.1) then
        print *,"Error: switch_elm_heat_flux only accepts 0 or 1 as inputs. Terminating div1d.exe."
        call exit
elseif(switch_elm_series.ne.0.and.switch_elm_series.ne.1) then
        print *,"Error: switch_elm_series only accepts 0 or 1 as inputs. Terminating div1d.exe."
        call exit
endif

time_relative = time - elm_start_time*delta_t
time_mod = time_relative
elm_heat_load = 0d0
elm_density_change = 0d0
elm_count = 0

do while (time_mod.gt.0d0)

        if (switch_elm_heat_flux.eq.1) then
                time_integrated_quantity = elm_expelled_heat
                if (gaussian_elm.eq.1) then
                        call maxwell_boltzmann_elm(time_mod,time_integrated_quantity,quantity,dquantity)
                else
                        call triangular_elm(time_mod,time_integrated_quantity,quantity,dquantity)
                endif
                elm_heat_load = elm_heat_load + quantity
                if (isnan(elm_heat_load)) then
                        elm_heat_load = 0d0
                endif
        endif

        if (switch_elm_density.eq.1) then
                time_integrated_quantity = elm_expelled_particles
                if (gaussian_elm.eq.1) then
                        call maxwell_boltzmann_elm(time_mod,time_integrated_quantity,quantity,dquantity)
                else
                        call triangular_elm(time_mod,time_integrated_quantity,quantity,dquantity)
                endif
                elm_density_change = elm_density_change + dquantity
                if (isnan(elm_density_change)) then
                        elm_density_change = 0d0
                endif
        endif

        elm_count = elm_count + 1
        time_mod = (time_relative - elm_count*elm_time_between*delta_t)*switch_elm_series

end do

end subroutine simulate_elm
```

Box B.5. Including boundary conditions in right_hand_side

The simulate_elm subroutine is called in the right_hand_side subroutine—responsible for calculating ydot, which contains the time derivatives of the values contained in the solution vector y—directly after it has called the subroutines responsible for the fluxes and sources.

```
call simulate_elm(elm_heat_load, elm_density_change, time)
```

The ramp rate due to the ELM at the current time is outputted as elm_density_change and then included as a boundary condition in the same was that the existing (user-defined input) density_ramp_rate had been included.

```
ydot(1) = density_ramp_rate + elm_density_change
```

The heat flux is added in the line that sets the boundary condition for the heat flux. It is added along with the contribution of the steady-state heat flux q_parX (which again previously existed).

```
ydot(2*Nx+1) = ydot(2*Nx+1) - (q_parallel(1)-(q_parX+elm_heat_load))/delta_xcb(1)
```