Ignition and Burn in Perturbed Inertial Confinement Fusion Hotspots

Jon King Tong

April 2019

Supervised by Professor Jeremy Chittenden

Declaration

I hereby certify that the material in this thesis, which I now submit for the award of Doctor of Philosophy, is entirely my own work unless otherwise cited or acknowledged within the body of text.

Jon King Tong
9th April 2019

Copyright

Abstract

This work explores how perturbations affect the performance of inertial confinement fusion hotspots through the inhibition of ignition and the effect on burn propagation post-ignition. To model this, a 3D Monte-Carlo charged particle transport module has been developed for the radiation-magnetohydrodynamics code Chimera.

The behaviour of the hotspot and the hotspot power balance in three alpha-heating regimes — self-heating, robust ignition and propagating burn — are explored in 1D through hydrodynamic scaling with capsule size and laser energy, demonstrating strong alpha-heating effects on the hydrodynamic evolution. The hotspot's definition affects calculated hotspot parameters, particularly in weak alpha-heating regimes where the boundary between the hotspot and the shell is less well-defined.

The impact of perturbations on the hotspot was explored using idealised spike and bubble perturbations. Less efficient PdV heating of the hotspot reduces the strength of the alpha-heating bootstrap. The spike significantly increased heat flow out of the hotspot from thermal conduction and alpha-heating; due to larger temperature gradients and surface area around the spike, and the closer proximity of cold, dense material to the primary fusion regions respectively. The bubble's effect on the hotspot power losses was minimal, but re-expansion into the bubble reduced confinement and truncated burn.

3D implosions based on National Ignition Facility Highfoot and High-Density Carbon (HDC) designs were perturbed using short-wavelength multi-mode and long-wavelength radiation asymmetry perturbations and scaled with capsule size and laser energy. The multi-mode yield increased faster with scale factor due to more synchronous PdV compression producing higher temperatures, and therefore stronger alpha-heating bootstrapping. Significant yield degradation for the Highfoot design resulted in only modest improvements in yield from scaling, whereas the less degraded HDC design demonstrated non-linear yield scaling. Perturbed implosions in the propagating burn regime exhibited fire-polishing due to thermal conduction and alpha-heating, in addition to "aneurysm"-like loss of confinement.

Role of the Author

The work in this thesis was done as a part of the Centre for Inertial Studies within the Plasma Physics group at Imperial College London. Chimera, sometimes known as Gorgon in its MHD form, has been in development by the team over a number of years. Many contributors have built different aspects of the code, including but not limited to the hydrodynamics, the radiation transport and the thermal conduction. The author was responsible for the development of the Monte-Carlo alpha-transport module within the code, the slowing-scheme of which was based on work by Mark Sherlock.

In addition, Kris McGlinchey's work on modelling National Ignition Facility experiments using Chimera was key to ensuring the accuracy of simulations on which this work was based. All other simulations and analysis were performed by the author.

Acknowledgements

First and foremost, I would like to thank my supervisor Professor Jeremy Chittenden for his invaluable support and guidance over the last three and a half years. Without his relaxed but steady approach and keen insight, my PhD experience would have been a very different story. A huge thanks must go out to the rest of the team — Kris, Chris, Aidan and Brian — for helping to proof-read and for many insightful (and some banal) conversations.

Thanks to all my office-mates — past and present — for some great company, and for putting up with my hoard of sports kit. Props to the Plasma Physics running intervals crew for being a staple part of my week, come rain or shine. Thanks also to the rest of the MAGPIE and Plasma Physics group, who have maintained a laidback yet intellectually stimulating atmosphere in which to work.

A huge thanks to my family, who have always supported me unconditionally. To my friends who have been there since the start of this journey, thanks for all the fun times. Finally, thanks to the ray of sunshine for brightening my life and supporting me through the difficult times.

Contents

List of Tables

List of Figures

1 Introduction

1.1 Fusion Energy

As the impact of global warming and climate change becomes more evident, the global pressure on governments to utilise renewable, carbon-neutral energy sources is rising, as evidenced by the recent Paris climate accord [10]. However, conventional renewables such as solar or wind power struggle to attain the scalability and reliability that has made traditional fossil-fuel-based power so prevalent, due to the variability and weather-dependence of the power output. Energy storage research is ongoing to find methods to compensate for the variability of solar and wind, but has many challenges in finding an appropriate medium that is able to satisfy the numerous requirements in capacity, stability, charge and discharge rates, cost-effectiveness, reliability and scalability [11]. The ecological and environmental impacts of the large land areas required for solar and wind farms also need to be taken into consideration. Alternatives such as hydroelectric or geothermal provide more reliable power output, but remain very specific in their location requirements. Nuclear fission, being neither wholly renewable nor fossil-fuel, falls in an unusual middle-ground between the two. Scalable, stable and reliable, it is nonetheless controversial, due to the extremely long-lived radioactive waste and the possibility of a nuclear disaster. Strong fear of radioactivity among the public, combined with disasters such as the 2011 Fukushima incident have resulted in stronger opposition to the use of nuclear fission as a major energy source. Indeed, following the Fukushima incident, Germany moved to phase out nuclear power entirely, although some countries such as China and India continue to expand the sector.

In light of this, nuclear fusion energy presents a highly attractive potential solution, combining key advantages of both renewables and non-renewables. In combining a near limitless supply of fuel that can be extracted from seawater, a lack of greenhouse gas emissions and a capacity for stable, high-density power generation, nuclear fusion offers an almost ideal power source. There is no risk of a nuclear disaster as can happen with runaway chain reactions in fission, due to stringent conditions required in the reactor to maintain the fusion reaction. Even relatively minor deviations from the optimal conditions can quench the reaction and halt any power output, and so fusion energy is inherently very safe — although naturally this raises the difficulty of studying fusion reactions, and therefore our ability to control and exploit it as an energy source.

It is perhaps worth noting that in recent years, the landscape of the fusion research industry is beginning to change. Growing interest from private enterprises into the potential of fusion energy has resulted in the emergence of a number of fusion-related

companies. The approaches taken by these companies vary greatly, with some similar to mainstream industry research, and some less so. While much work needs to be done to validate each of these approaches, the increased diversity nonetheless highlights a changing dynamic in the field.

Differences in binding energies between the products and the reactants in a nuclear reaction manifest as mass differences, and thus can be used to generate energy (via $E = mc^2$). This is the core concept on which nuclear fission and fusion energy are based. Fusion reactions generally require extremely high temperatures in order to overcome the Coulomb repulsion between the colliding nuclei. While these conditions are abundant in astrophysical circumstances, there are limits to the achievable temperatures and pressures in a man-made reactor vessel. Because of this, the deuterium-tritium reaction $D + T \rightarrow {}^4He + n$ is commonly considered to be the most viable for energy generation, due to a significantly greater reaction cross-section at the conditions accessible in the lab (where the cross-section indicates the probability of a reaction occurring upon collision).

The simplest approach of a beam-target setup results mostly in heating of the target due to Coulomb scattering, rather than fusion. It follows that beam fusion is not a viable approach to achieving fusion, unlike thermonuclear fusion – where the fuel is heated to (and contained as) a plasma at very high temperatures. The two primary methods of confining the fusion plasma are with magnetic fields (magnetic confinement fusion, MCF) and with the plasma's own inertia (inertial confinement fusion, ICF).

MCF approaches generally aim to confine a low density, high-volume plasma continuously over an extended period of time, a steady-state process in which the nuclear fusion sustains the plasma and provides the heat from which electricity can be generated. The most commonly used device is the tokamak, which confines the plasma in a toroidal configuration. Many organisations engaging in fusion research around the world focus on this approach, including the JET and MAST at the Culham Centre for Fusion Energy (CCFE), and the ITER collaboration between the EU, USA, China, Japan, South Korea, India and Russia, based in the south of France.

ICF approaches are generally transient, using the plasma's own inertia to confine a very small volume of dense plasma over a short time period, during which fusion reactions can occur.

1.2 Inertial Confinement Fusion

In this section, we explore the physical basis of inertial fusion energy, and the different approaches taken by the field to achieve it.

1.2.1 Hotspot Ignition

Inertial confinement fusion (ICF) seeks to achieve fusion by using inertia (due to fuel's own mass) in order to confine the fuel. In contrast to MCF's steady-state confinement, ICF is an inherently pulsed process, and needs a significant proportion of the fuel to be

burnt in fusion in order to produce power efficiently. The burn efficiency characterises the fraction of the fuel that is burnt in fusion, and is defined as $\Phi = N_{fus}/N_0$ for total fusion reactions, N_{fus} and number of pairs of DT particles contained in the fuel, N_0. For equimolar DT, this can be found to be [12, 13]:

$$\Phi(t) = \frac{t/2\tau_r}{1 + t/2\tau_r} \tag{1.1}$$

where $\tau_r = (n\langle\sigma\nu\rangle)^{-1}$ for total number density n and reactivity $\langle\sigma\nu\rangle$.

For a plasma sphere of initial radius, R_0, mass, m_0, and volume, V_0, expanding in a rarefaction wave at sound speed, c_s, the radius of the non-rarefied region $R = R_0 - c_s t$. Assuming that only the non-rarefied region, $V(t)$ undergoes fusion reactions, the effective confinement time is less than the time taken for the entire sphere to disassemble, $\tau_{conf} = R_0/c_s$ and is given by:

$$\int_0^{\tau_{conf}} \frac{V(t)}{V_0} dt = \int_0^{\tau_{conf}} \frac{R^3}{R_0^3} dt = \int_0^{\tau_{conf}} \left(1 - \frac{c_s t}{R_0}\right)^3 dt = \frac{R_0}{4c_s} \tag{1.2}$$

Substituting this into equation 1.1, using $n = \rho/m_0$ and rearranging gives:

$$\Phi \approx \frac{\rho R_0}{H_B + \rho R_0} \tag{1.3}$$

where $H_B = 8c_s m_0/\langle\sigma v\rangle$ is the burn parameter. Taking reference values of $H_B = 70 kg/m^2$ [7] and a burn efficiency of 30%, the confinement parameter ρR required is found to be $\sim 30 kg/m^2$.

In order to minimise damage to reactor chambers, the energy released from each fuel capsule must be limited to a few GJ, and therefore the fuel mass restricted to $O(10^1) mg$. Considering the compression ratios required, with the density of solid DT, $\rho_{DT} = 225 kg/m^3$, then for a fuel mass of $1 mg$ the required density works out to $3 \times 10^5 kg/m^3$, i.e. a compression ratio of ~ 1300 [7]. This is easiest to achieve in a convergent geometry, and so the fuel capsule is usually imploded and compressed spherically using high-powered radiation.

An example fuel capsule set-up is shown in Fig. 1.1 for a spherical NIF capsule, with distinct ablator and solid DT layers outside the central region containing gaseous DT. The ablator layer is often doped with a material of high atomic number (high-Z) like Germanium or Silicon to prevent hard x-rays from penetrating too deep and preheating the capsule, which reduces its compressibility and increases the energy requirements for a given compression. The dopant concentrations are graded in several layers in order to reduce hydrodynamic instability growth. The fabrication of the solid DT layer is time-consuming since it must be cryogenically cooled to 18K, in order to freeze the DT.

While isentropic compression would minimise the entropy gain and energy required for compression to a given ratio, the compression must also be fast in order to generate the velocities required to attain the temperatures needed in the hotspot. The strong shocks
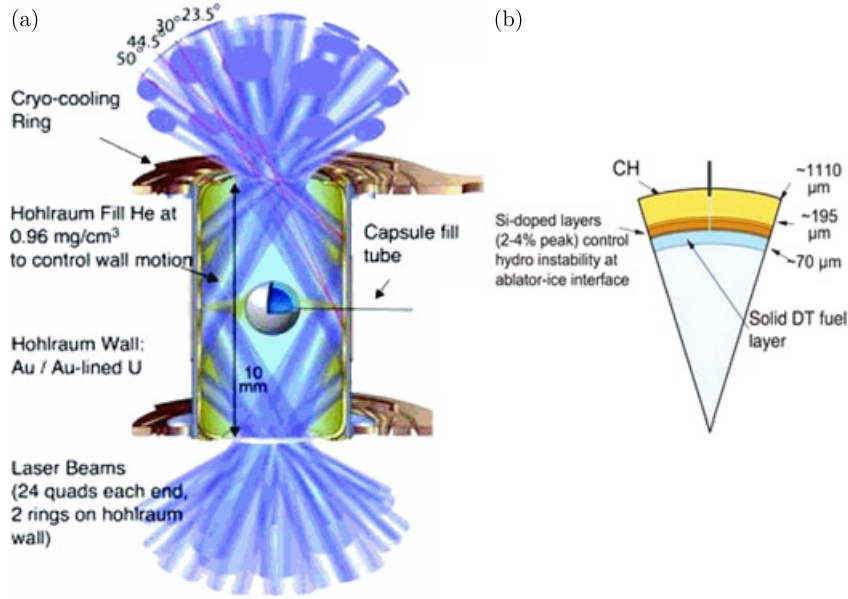
Figure 1.1: (a) An example NIF capsule target setup in a hohlraum and (b) an example NIF target of around ~ 1mm in radius, with the ablator dopant graded in multiple layers in order to mitigate hydrodynamic instabilities. Reprinted from [1] with the permission of AIP Publishing.

driven by the radiation pulse produce significant entropy, and the compression across each shock front is limited to factors of $\sim 4 - 6$. However, the use of multiple shocks allows for both fast and near isentropic compression (with ideal isentropic compression occurring in the limit of infinitesimally strong shocks of infinite number).

The implosion is propelled in a manner similar to an ablative rocket, using a driver to heat up and expand the ablator layer on the outside of the capsule. As it expands, conservation of momentum drives the capsule inwards, following the rocket equation:

$$v_{shell} = \frac{P_a}{\dot{m}} \ln \left(\frac{m_0}{m_f} \right) \equiv v_{exhaust} \ln \left(\frac{m_0}{m_f} \right) \tag{1.4}$$

for an ablation pressure, P_a, a mass ablation rate per unit area, $\dot{m}_a$, an initial shell mass, m_0 and a final shell mass, m_f. Typically, ablation-driven rockets are much less efficient than ideal rockets because the driver also heats the exhaust as it drives the implosion, resulting in an implosion efficiency of between 5-20%, depending on the driver type. The total work done during the implosion is $P_a V$ for a shell enclosing an initial volume, V, where the final volume $V_f \ll V$. Thinner shells will produce a faster implosion velocity than a thicker shell for the same drive due to the lower shell mass, but also risk greater disruption from hydrodynamic instabilities.

For a fuel capsule driven with driver energy, E_d, which produces total fusion energy, E_{fus}, the energy gain is defined as:

$$G = \frac{E_{fus}}{E_d} \tag{1.5}$$
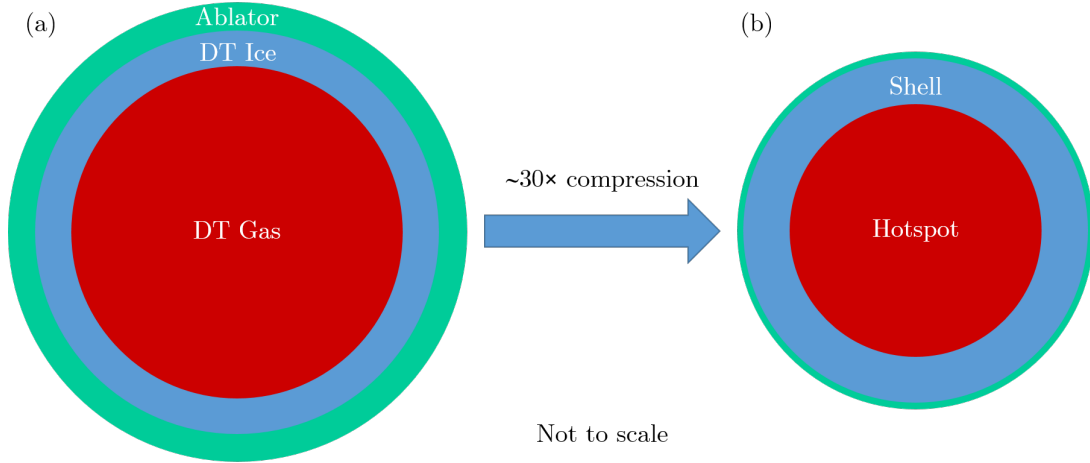
Figure 1.2: A schematic showing: (a) the general structure of the capsule, consisting of layers of (from outermost to innermost) the ablator layer (green), DT ice (blue) and DT gas (red); and (b) the structure of the capsule at peak compression, consisting of a central high-temperature, low-density hotspot (red), an outer layer of low-temperature, high-density DT fuel (blue) and the remnants of the ablator (green).

G needs to be large enough to account for energy-conversion inefficiencies in the driver, and the thermal cycle to convert the energy into electricity. Assuming a thermodynamic cycle efficiency of 40%, and requiring 25% of the electrical energy to be used to power the reactor, then a driver efficiency of 10-30% gives a gain requirement of ∼30-100. In addition, considering the efficiency of coupling driver energy into fuel energy, η_{imp}, the fuel burn-up Φ and the internal energy of the plasma, $3k_BT$, this gives:

$$G = \frac{17.6MeV}{2 \times 3k_bT} \times \Phi \times \eta_{imp} \tag{1.6}$$

for two reacting ions in the plasma, where $17.6MeV$ of energy is released by each DT reaction, and we consider the internal energy of both the electrons and ions of each reactant. Taking relatively optimistic values of $\Phi \sim 0.3$ and $\eta_{imp} \sim 0.1$ gives $G \sim 20$. This therefore lowers the gain below the requisite levels if heating the entire fuel capsule to ignition temperatures, known as volume ignition.

Hotspot ignition circumvents this problem, by driving only a small portion of the whole fuel mass (called the hotspot) to ignition temperatures, and using this to initiate burn throughout the rest of the capsule. Figure 1.2 shows (a) the general pre-implosion structure of the capsule, with layers of ablator, DT ice and DT gas, and (b) the peak compression structure of the capsule at the final stage of the implosion. This consists of two distinct regions – <5% of the total fuel mass within the high-temperature, low-density central hotspot and the rest in the low-temperature, high-density surrounding shell of compressed fuel and remnants of the ablator layer. The hotspot needs to not only be self-heating but produce enough fusion to propagate the burn throughout the rest of the shell, which needs to absorb the majority of the energy contained within the fusion alpha-particles.

A typical hotspot ignition implosion involves multiple phases [7]. During the initial "ablation" phase, the radiation drives an ablation front inwards as the outer surface of the capsule is ablated. Ahead of the ablation front, a shock accelerates and compresses the shell as it propagates inwards, driven by the ablation pressure. The initial shock is driven at a relatively low ablation pressure in order to minimise the preheat of the shell. The "acceleration" phase entails a series of jumps in the radiation drive temperature, along with a corresponding increase in ablation pressure. Each jump drives a successively stronger shock into the capsule, calculated to coalesce just inside the ice layer of the capsule, aiming to create fast and near-isentropic compression through a series of shocks as described earlier. The radiation-drive is switched off after the majority of the driver energy has been deposited, leaving the capsule material to freefall inwards, where it is decelerated by rebounding shocks from the centre in the "deceleration" or "stagnation" phase. This deceleration leaves relatively low-velocity material stagnating in the centre of the capsule, which heats up as it is compressed by the rest of the dense fuel. The innermost material forms the "hotspot", which continues to be heated by mechanical work up until peak compression, when the rest of the confining shell has stagnated (and the capsule is at its smallest). Fusion reaction product energy deposition further heats the hotspot, facilitating more fusion in a positive feedback or "bootstrapping" process. If this is strong enough, then the hotspot will ignite and the heat transport from the hotspot (via electrons and/or fusion products) propagates a burn wave outward throughout the rest of the dense shell during a "burn" phase, igniting the rest of the capsule. This will be further elaborated on in section 2.4 and chapter 4.

1.2.2 Approaches to Hotspot Ignition

The two approaches to achieving hotspot ignition in ICF most commonly used are direct- and indirect-drive, which refer to the use of high-powered radiation to either directly or indirectly drive a spherical implosion.

Direct drive is so-called as the (typically) laser radiation illuminates the capsule surface directly, where the electrons are heated and electron conduction transfers energy to the ablation front. However, long-wavelength non-uniformities can be generated depending on the beam placement and the total number of beams, and short-wavelength non-uniformities due to the intensity profiles of individual beams. In addition, laser-plasma instabilities can affect the energy-coupling of the laser to the capsule, as well as accelerate fast electrons which then preheat the fuel. Facilities conducting direct-drive experiments include the OMEGA laser [14] at the University of Rochester Laboratory for Laser Energetics (LLE), the Nike laser [15] at the Naval Research Laboratory (NRL) and the GEKKO XII laser [16] at the University of Osaka. Although the National Ignition Facility (NIF) at the Lawrence Livermore National Laboratory (LLNL) is primarily an indirect drive facility, polar direct drive [17] experiments have also been recently fielded on NIF [18].

Indirect drive uses high atomic number (Z) enclosures called "hohlraums" around the

fuel capsule to generate an x-ray bath, which ablate the capsule surface and drive the implosion. Although the laser-induced non-uniformity effects will also be relevant for the laser deposition into the hohlraum, the use of the hohlraum reduces the impact of these effects on the capsule. As a result, indirect drive allows for more uniform intensity across the capsule surface. In addition, the x-rays generated by the hohlraum are absorbed directly near the ablation front, whereas laser energy is deposited into electrons at the critical surface. The energy is then transported to the ablation front through thermal conduction, a process which becomes less efficient as the distance between the ablation front and the critical surface increases. As such, the more efficient absorption of radiation in indirect drive therefore produces higher ablation pressure and velocity. However, the overall efficiency of coupling from laser energy to shell kinetic energy is lower for indirect drive. This is predominantly due to the difference in surface area between the hohlraum and the capsule — such that much less energy is incident on the capsule surface than was incident on the hohlraum surface — with a minor effect from the conversion of the laser into x-rays. NIF is used primarily for indirect drive experiments, and the next section (Section 1.3) will go into further detail on the indirect drive approach and the difficulties involved in achieving ignition on NIF. The Laser Mega Joule (LMJ) facility in France is based on a similar design to NIF, and thus also plans primarily to engage in indirect drive experiments.

1.2.3 Alternative ignition schemes

Alternative approaches to ICF often involve the separation of the compression from the heating and ignition of the fuel, as in shock ignition and fast ignition, or the use of magnetic fields. Separating the heating and ignition phases allows the use of slower implosion velocities to compress the capsule, which reduces disruption due to hydrodynamic instabilities. Lower laser intensities also may also reduce the impact of laser-plasma instabilities. However, the resultant lower temperatures and pressures require an additional source of energy to trigger ignition.

In fast ignition, energetic particles such as electrons or ions are used as an external heating source [19, 20]. These particles are typically generated by a high-intensity ignitor laser pulse, with different methods of accessing the compressed plasma. The first [19] involves using a laser to create a channel in the coronal blow-off plasma, through which a second, high-intensity laser propagates in order to generate MeV electrons, in a process known as hole-boring [21]. While this first uses conventional spherical capsules, the second approach incorporates a cone into the capsule geometry [22]. This cone avoids the need for a channelling pulse and shines the high-intensity, particle-generation pulse directly onto the tip of the cone, where the electrons are generated very close to the imploded plasma.

Shock ignition uses a spike in the laser intensity at the end of the laser pulse to propagate a second, strong, spherically convergent shockwave into the centre of the capsule [23, 24]. This second shock collides with the initial compression shock rebounding from

the centre of the capsule, and results in higher hotspot stagnation pressures than in conventional hotspot ignition. This scheme typically requires direct drive in order to produce the sharp change in drive intensity required to generate the final shock. This is due to the heat capacity of the hohlraum, which significantly slows the rise time in laser energy to give a much slower rise in radiation temperature. In addition, the large ratio between the hohlraum radius to the capsule radius (known as the case-to-capsule ratio, or CCR) at the time of ignitor shock launch would reduce the coupling efficiency and result in a much weaker shock [25].

Magnetised ICF [26, 27] aims to use magnetic fields embedded in the target as a way of reducing the heat losses within the hotspot, while magnetised liner inertial fusion (MagLIF) combines laser heating with this using an axial magnetic field in a cylindrical target [28, 29].

In addition, interest in a volume ignition approach originating with double shells [30, 31] has been rekindled with Molvig's Revolver triple shell design [32]. The core concept involves the use of a high-Z pusher shell, which reduces the convergence required to achieve a given shell density due to the higher density of the material. By contrast, hotspot ignition requires a DT shell in order to increase gain through burn propagation. The triple shell design can be considered as a double shell variant with improved "impedance matching" between the shells to maximise efficiency of kinetic energy conversion. The high-Z shell is highly effective at capturing the radiation emitted by the fuel, such that the re-emitted radiation can be used to heat the fuel to ignition temperatures before stagnation.

1.3 National Ignition Facility

The National Ignition Facility (NIF) is a laser facility at the Lawrence Livermore National Laboratory (LLNL), with a 192-beam laser capable of delivering 1.8MJ with a maximum power of 500TW at 351nm in a 10m-diameter target chamber. It was designed with the ignition and thermonuclear burn of deuterium-tritium (DT) capsules in mind, and most of the NIF ICF experiments are based around the indirect-drive approach described above. The physical basis for using NIF to achieve indirect-drive ignition is documented by Lindl *et al.* [6, 33]. This section briefly outlines some of the experimental and modelling progress made over the past decade at NIF.

1.3.1 National Ignition Campaign

The initial series of experiments in the National Ignition Campaign (NIC) ran between 2009-2012, and had the objectives of 1) achieving ignition and 2) demonstrating a variety of diagnostic and experimental platforms and techniques useful for future HEDP and ICF applications. The "Lowfoot" implosion design aimed to minimise the energy threshold for ignition by maximising fuel ρR, with a target of attaining high gain. This design used a CH ablator setup, similar to the one shown in figure 1.1.

Although the campaign did not manage to achieve ignition, progress was made on scientific understanding, technical expertise and the development of an array of diagnostics and infrastructure. Shock-timing experiments showed adequate control to obtain specified adiabats, and the target manufacture and characterisation precision greatly improved [34]. NIC also revealed a number of performance-degrading factors on implosions which lowered the yield below initial expectations. Foremost among these is hydrodynamic instability growth stemming from a variety of sources, including but not limited to: the contact between the tent membrane and the capsule surface, the indent from the fill-tube, and asymmetries in the radiation drive.

These perturbation seeds grow due to the Rayleigh-Taylor (RT) and Richtmeyer-Meshkov (RM) instabilities on the ablation front, the ablator-ice surface (during the ablation/acceleration phases) and on the ice-gas surface (during the deceleration phase). The initial perturbations can proceed to seed instabilities later in the implosion as the growth "feeds through" from the outer (ablator-ice) to the inner (ice-gas) surfaces, with a more pronounced effect for longer-wavelength perturbations. One can distinguish between two types of perturbation, "shape" and "mix", which generally refer to long and short wavelength Legendre l-modes (with $l = 2$ or 4 and $l \gtrsim 50$ respectively). The "mix" instabilities allow material from the ablator to penetrate and mix into the hotspot, with the high-Z material increasing the radiative losses from the hotspot, while the "shape" perturbations inhibit the conversion of implosion kinetic energy to thermal energy. This is due to the non-uniformity of deceleration between different parts of the capsule, such that some regions are decelerating before others [35].

1.3.2 Highfoot

After the end of NIC, various aspects of the implosion were modified in order to make progress towards achieving ignition, as in the Highfoot campaign [36, 37]. The Highfoot campaign used a differently shaped laser-pulse to drive a 3-shock implosion with a higher adiabat and therefore a lower convergence ratio [38], reducing susceptibility to instability growth at the ablation front. By increasing the initial picket (that is, the first part of the pulse), the first shock driven into the capsule was stronger and faster, and the higher trough increased the ablation velocity (and therefore ablation-stabilisation (see Section 2.3). Doing so reduced the maximum possible gain of the implosion design, but allowed greater control over the implosion and significantly reduced RT instability growth at both high and low modes. Milestone performance was achieved in which the alpha-particle energy deposition doubled the overall yield of the implosion [39], reaching stagnation pressures of $> 200 Gbar$ and hotspot temperatures of $\sim 5 keV$.

Extensive computational modelling in both 2D [40] and 3D [2] allowed further investigation into the mechanisms of performance degradation. 2D integrated simulations of the hohlraum and capsule using LLNL's radiation-hydrodynamics code HYDRA [41] investigated the effect of low-mode asymmetries [40, 42]. These low-mode asymmetries showed a significant impact on the yield, characterised by the residual kinetic energy

(RKE) of mass flows within the hotspot due to imploding fuel which did not decelerate and stagnate properly. Strong correlations were found between yield and RKE, with the level of alpha-heating increasing significantly towards 1D-like performance as RKE $\rightarrow$ 0. The source of these low-mode asymmetries is the radiation drive; effects such as cross-beam energy transfer (CBET) [43] and laser-plasma instabilities (LPI) [44] can alter the relative power between the inner and outer laser beams incident on the hohlraum near the central circumference (the waist) and near the laser entrance hole (LEH) respectively. CBET can increase the equatorial drive from the inners relative to the polar drive from the outers earlier on in the implosion. In addition, the incidence of the outer beams on the hohlraums can ablate high-Z plasma, which can affect further laser propagation to the wall as well as to the waist [45]. This reduced laser power at the waist can result in a stronger drive at the poles of the capsule, pushing the capsule from a prolate spheroid shape earlier on towards an oblate spheroid. This effect of a varying P_2-mode over time is known as "swing".

3D capsule-only simulations [2] revealed the extent to which 3D effects can perturb the implosion from an ideal 1D spherical symmetry for both Lowfoot and Highfoot implosions. These simulations incorporate effects of the capsule fill-tube, support tent [46, 47] and surface defects [48], in addition to the aforementioned radiation drive asymmetries. The simulations confirmed the reduction in ablation front instability growth from Lowfoot to Highfoot, and again identified the low-mode asymmetries as the next-most-significant factor affecting implosion performance. Although the tent-driven perturbations were reduced, they remained significant, particularly for the strongly driven implosions with implosion velocities of $> 370 km s^{-1}$.

This particular study highlights the capability of simulations to recreate experimental data, with agreement across several measures of performance (e.g. neutron yield, DSR, bang time) for implosions spanning different regions of parameter space (e.g. in compression ratio, ρR and yield). However, in spite of this, discrepancies remain between simulation and experimental diagnostic data. For example, the burn-averaged ion temperatures are consistently lower in simulations. Potential sources of error are wide-ranging, such as from the physical models used in the simulation for e.g. thermal conductivity [3], or the exclusion of other physics such as self-generated magnetic fields [49].

1.3.3 High-Density Carbon

Following the Highfoot campaign, the experimental direction moved further towards more robust and predictable designs, with high-density carbon (HDC) (i.e. nanocrystalline diamond) being explored as an ablator. Using a HDC ablator allowed for capsules with thinner shells and larger inner radii than for an equivalent outer radius CH capsule, corresponding to larger DT shell and hotspot masses. HDC is a more efficient ablator than CH as the ablation front forms at a larger radius from the thinner shell, and therefore has a larger surface area and absorbs more energy, allowing for similar gains on
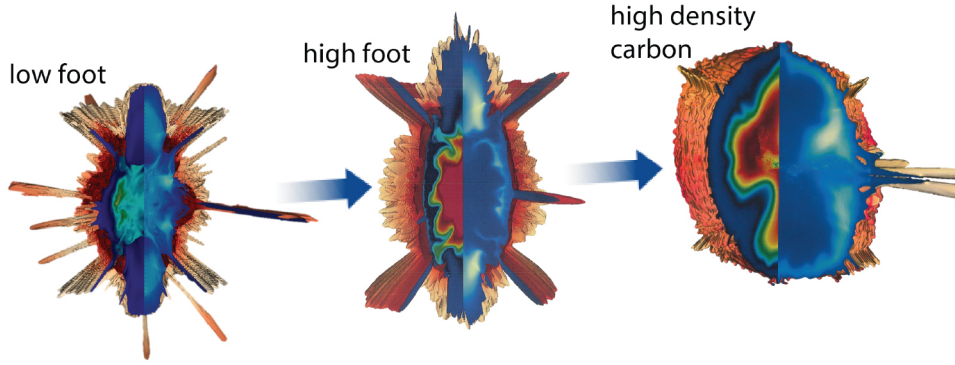
Figure 1.3: Capsule implosion progress between the Lowfoot, Highfoot and HDC campaigns shown at their respective bangtimes using post-shot 3D HYDRA simulations of NIF implosions N120405, N140819 and N161023. The left-half displays ion temperature, and the right-half shows the density. Reproduced from [2] and [3] with the permissions of AIP Publishing and IAEA .

higher adiabats, or conversely higher gains on similar adiabats [50]. The surface roughness of HDC can be polished to $10\times$ better than for CH, since it is a nanocrystalline structure. The higher density also allowed for shorter laser pulses to be used ($8 - 10ns$ compared to $15 - 22ns$ for CH), and by extension the use of near-vacuum hohlraums [51]; the gas-fill is primarily to prevent the plasma ablated off the hohlraum itself from filling the hohlraum during long pulses [52]. Near-vacuum hohlraums demonstrated better efficiency due to reduced backscatter, less fast-electron preheat and better hohlraum-capsule energy coupling, in addition to better symmetry control from avoiding the use of CBET for symmetry control [53]. However, the HDC ablation front scale lengths are smaller, and if the HDC does not melt during compression, the inherent crystalline grain structure can provide seeds for hydrodynamic instability growth [54]. The initial shock temperature must therefore be high enough to melt the HDC in order to avoid this structure imprinting on the shock, which necessarily increases the adiabat of the capsule.

The resulting experimental campaign demonstrated significant improvements in symmetry control [55], particularly regarding low-mode asymmetries. As a result, the HDC campaign produced landmark yields of 1.5×10^{16} and 1.7×10^{16} with shots N170601 and N170826 [9]. Figure 1.3 (reproduced from [2, 3]) highlights the progress in implosion symmetry between the Lowfoot, Highfoot and HDC campaigns. The Lowfoot hotspot is not very well formed, while the HDC hotspot is significantly more spherical in shape than the Highfoot hotspot. However, yield degradation due to the fill-tube in these experiments is more prominent [3], having successfully reduced the impact of the tent and radiation asymmetry perturbation sources. The fill-tube yield degradation occurs both due to the entrainment of ablator and fill-tube material into a jet through the hotspot [3], and the shadows cast by the fill-tube on the capsule surface during the initial X-ray drive [56] seeding instability growth. Reducing the fill-tube diameter was a factor in producing the landmark yields above [57].

1.3.4 Exploring the design space

The number of variables in implosion design which can be modified is extensive, including but not limited to the case-to-capsule ratio [58], the hohlraum gas-fill density [59], the laser pulse shape and length, the ablator material [60, 61], the hohlraum material [62], and even the hohlraum shape [63, 64]. There has also been some investigation into counterbalancing the anticipated asymmetries by adjusting the initial shape of the capsule (known as "shimming") [65]. Some of these, such as different ablator materials, have already been mentioned in the experimental progression towards the current HDC designs.

Further work on CH ablator capsules focuses on extending the pulse-shape modifications made from Lowfoot to Highfoot [66]. While the Highfoot pulse shape drove a more hydrodynamically stable implosion, it was also less compressed. Adiabat-shaping aims to maintain the implosion stability of the Highfoot pulse shape in addition to achieving the higher compressions as seen in the Lowfoot pulses. This is achieved by lowering the laser power of the "trough" between the first and second pulses, such that the strong first shock decays in the ablator and enters the capsule as a weaker shock, setting the fuel on a lower adiabat while keeping the ablation front on a high adiabat [67, 68]. Experiments have been done with both 3-shock [69] and 4-shock [70] laser-pulses, and have demonstrated significantly increased fuel compression while maintaining reduced ablation-front instability growth. The experiments demonstrated reduced instability growth relative to Lowfoot, and improved compression relative to Highfoot [71]. However, fuel compression was reduced in a high-power 3-shock compared to the low-powered 3-shock, thought to be due to increased electron preheat [72].

Beryllium is another alternative ablator which is undergoing investigation due to its low opacity and high density. This results in greater ablation rates, pressures and velocities, and therefore greater ablative stabilisation to hydrodynamic instabilities [60, 61]. The drive temperature requirements are lower, allowing improvements to the x-ray drive symmetry through the use of larger case-to-capsule ratios. However, the increase in ablated plasma in the hohlraum may affect laser propagation, and the microstructure could seed instabilities [73]. Initial experiments with beryllium ablators were predominantly affected by low-mode radiation asymmetries [74, 75]. However, more recent experiments demonstrated improved predictability and control through to stagnation [76].

The ablators CH, HDC and beryllium each have their own advantages and disadvantages, both in regards to radiation drive characteristics and susceptibility to various perturbations. Simulation studies comparing the current and extrapolated performance of the ablators have been done, both integrated hohlraum-capsule [61] and more detailed capsule-only [60]. They indicated that more data are needed to better understand and characterise beryllium ablators. Although beryllium is expected to have the highest potential yield, it is also the most susceptible to both low-mode perturbations and surface roughness-induced short-wavelength mix. While the fill-tube currently affects HDC the most, the reduced $5\mu m$ fill-tube is expected to have a significantly smaller impact on

all ablators. The tent impacts current CH and HDC designs the most, but the extrapolated beryllium design is expected to be the most vulnerable. A key observation from the study is that the ablator choice is but one variable in the design space, and that other factors such as case-to-capsule ratio or pulse shape will also strongly influence the results. The set of parameters that are ideal for one ablator may not produce the best results for another ablator.

Independent of the choice of ablator, ways are being developed to mitigate and reduce the impact of various seeds of hydrodynamic instability growth. The capsule mounting tent is one such seed, and investigations are underway into alternative capsule mounts [77, 78]. These methods are aimed at either removing the tent completely, or reducing the tent's impact on the capsule hydrodynamics. The "tetra-cage" involves the use of four $2.5\mu m$-thick carbon nanotube wires to reduce the total contact area, with two parallel wires supporting the capsule from below, and a further pair at the top running perpendicular to the other pair. The cantilevered fill-tube uses the fill-tube to hold the capsule, and supports the fill-tube $100 - 300\mu m$ from the capsule surface using a $12\mu m$-thick SiC rod, while the "fishing-rod" instead supports the fill-tube by placing it inside another, thicker ($30\mu m$) tube at $300\mu m$ from the capsule surface. Polar tents instead aim to reduce the contact area between the tent and the capsule by attaching the tent further up the hohlraum, strengthening the tent with an $8nm$-thick layer of Carbon (facing away from the capsule). The "foam-shell" uses a $200\mu m$-thick layer of SiO_2 foam between the capsule surface and the tent to dampen the tent's impact. Initial hydrodynamic instability growth experiments show promising results, with all three tent-less methods showing a reduction in total area affected compared to the nominal tent, but an increase in shadows cast by x-rays. The foam shell prevented tent-related perturbations but demonstrated significant perturbations from target handling and manufacture, in addition to unexplained 3D broadband structures, while the polar tent was also dominated by unexplained 3D structures rather than the expected smaller radius tent perturbation. Overall, the investigation is far from conclusive and much work remains to be done [78].

The expansion of the high-Z plasma from the hohlraum wall has been seen to occur across a range of low- and intermediate gas-fills in hohlraums with a variety of ablators and hohlraum sizes [58]. As previously mentioned, this bubble of plasma can affect the propagation of the inner laser beams and therefore the overall symmetry of the radiation drive [45]. One method to address this involves changing the shape of the hohlraum by increasing the radius where the outer beams are incident to make an "I-raum", i.e. a hohlraum with a cross-section similar to that of an "I" [79]. By simply moving the wall further from the centre, the impact of the bubble is reduced as it takes longer to reach the inner beams. Other approaches to address this bubble involve increasing the gas-fill density or using a low-density foam.

1.4 Perturbations and igniting hotspots

It is clear that much work both has been and is being done to mitigate the impact of the numerous sources of perturbations which affect the performance of ICF implosions. Huge progress has been made over the past decade in identifying perturbation sources and minimising their effects on implosion performance. Experiments are now entering a regime in which alpha-heating is a significant contributor to the performance, reaching yields of $\sim 2 \times 10^{16}$ with a gain of ~ 2 relative to the kinetic energy of the shell [9]. However, implosions remain affected by perturbations, albeit to a significantly reduced degree.

This work focuses on the interactions between perturbations and hotspots undergoing ignition and burn. Much of the existing work in this area focuses on the growth rates of Rayleigh-Taylor Instabilities (RTI) in the deceleration phase. Lobatchev and Betti [80] established the significance of mass ablation from the shell's inner surface on the growth rates of deceleration phase RTI, using both analytic and 1D *LILAC* calculations. Atzeni demonstrated the ablative stabilisation effect of thermal conduction and alpha-heating from the hotspot on the RTI growth rates using the 2D *DUED* code [81] with single-group diffusive alpha-particle transport and bremsstrahlung radiative losses [4]. Notably, the impact of the perturbation on the hotspot formation was also considered, showing a delay in the ignition of the hotspot in addition to a reduction in size of the hotspot. Figure 1.4 shows the time-evolution of an ICF capsule perturbed by a large amplitude, $l = 16$ mode perturbation undergoing ignition and burn in this study, showing significant shell distortion and hotspot size reduction.

Schiavi and Atzeni [5] also explored the isolated perturbation growth rates for spikes and bubbles of varying extents, again using *DUED* in 2D with multigroup alpha-particle diffusion and bremsstrahlung losses. Here, better ablative stabilisation was found for spike perturbations compared to bubble perturbations, as demonstrated in figure 1.5 comparing perturbation growth for spike and bubble perturbations at three widths using a "classical" model with just hydrodynamics and bremsstrahlung losses, and a "fusion" model including thermal conduction, fusion reactions and alpha-particle transport. Fan *et al.* [82] also considered the growth rate of deceleration phase RTIs, finding that the increased mass ablation from direct deposition of alpha-particle energy into fuel in the dense shell layer was a stronger stabilising effect than the destabilisation of increased deceleration due to increased hotspot temperature and pressure, using the 2D code LARED-S with multi-group radiation and alpha-particle diffusion.

Levedahl and Lindl's [83] ignition threshold based on the implosion velocity and compressibility also included a method to incorporate the impact of cold material mix; this was done through an increase in velocity proportional to the ratio of the mixed (full) to the clean (reduced) hotspot radius. Kishony and Shvarts [84] built on this concept of a clean-volume analysis for perturbed hotspots by extending it to lower perturbation modes. This analysis was used in conjunction with a 2D self-similar analysis to develop a model for the ignition and gain of perturbed hotspots.
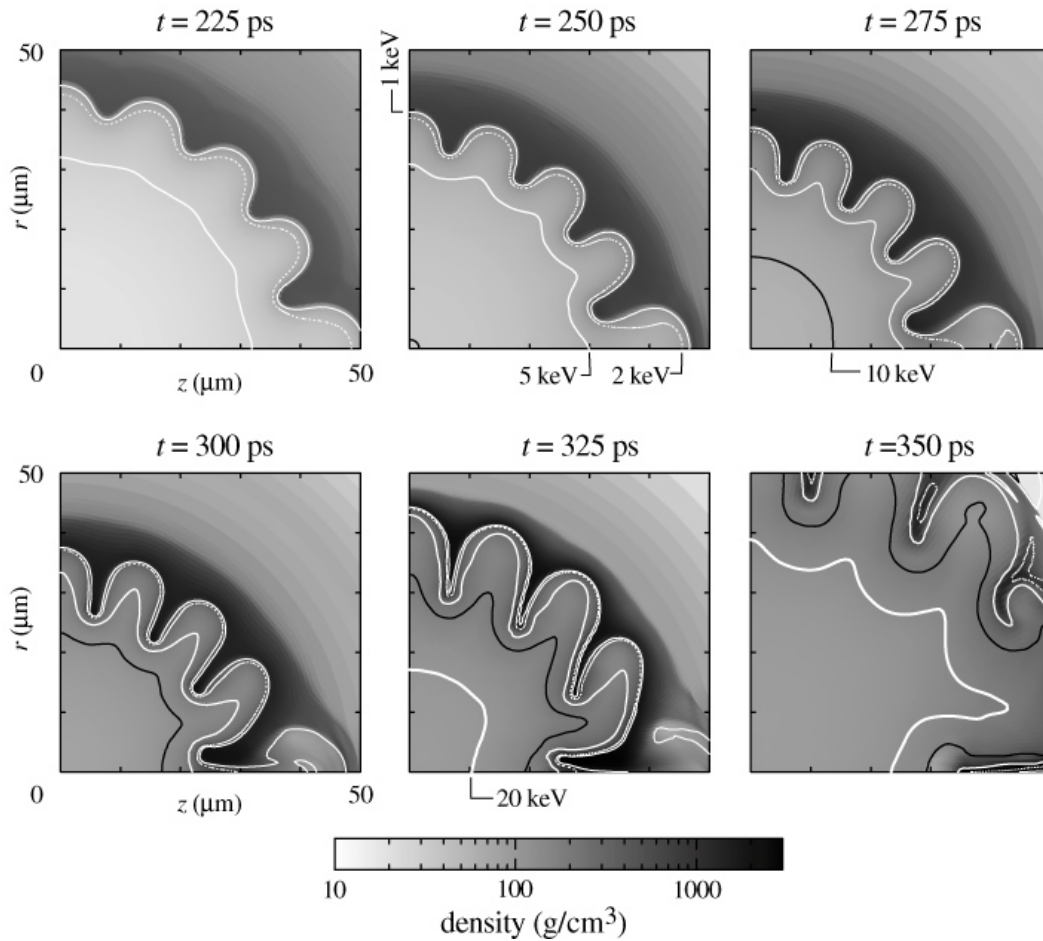
Figure 1.4: Time-evolution of a 2D *DUED* simulation of an ICF capsule perturbed by a large amplitude $l = 16$ perturbation undergoing ignition and burn, showing greyscale density maps and selected ion temperature contours. Reproduced from [4] with permission. © IOP Publishing. All rights reserved.
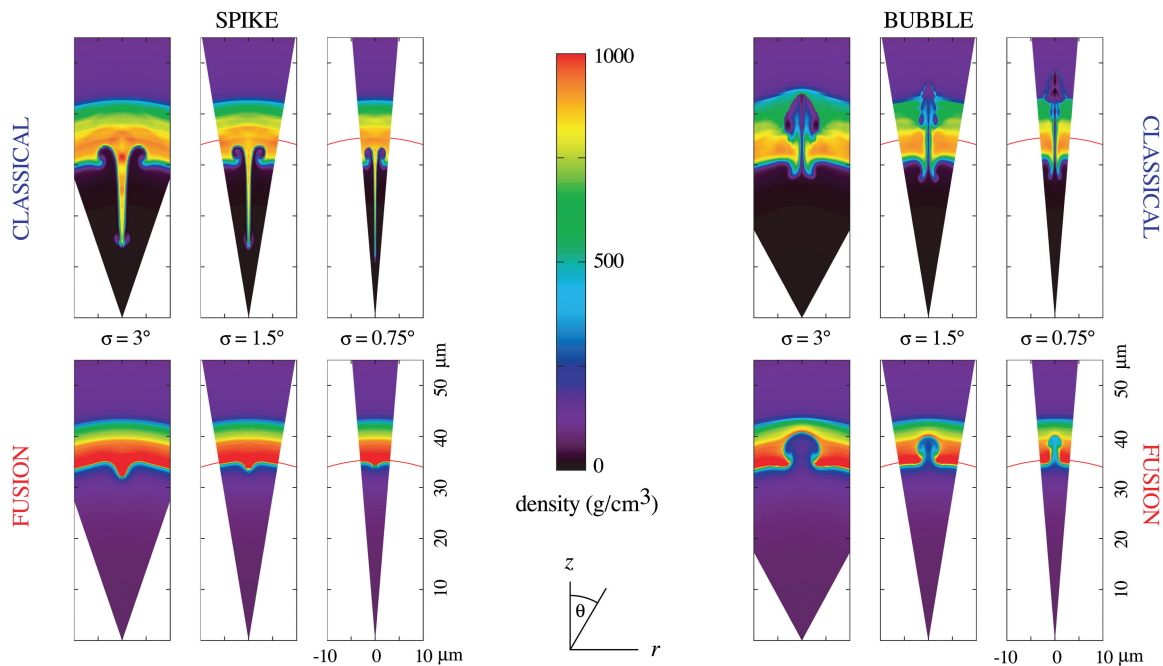
Figure 1.5: Comparison of the growth of single spike (left) and single bubble (right) perturbations using the "classical" and "fusion" models for three angular widths and identical initial amplitudes, performed using *DUED*. Reproduced from [5] with the permission of AIP Publishing.

More recent works include models for hotspot ignition under perturbed conditions by Springer *et al.* [85] and Cheng *et al.* [86]. Springer *et al.* [85] developed a 3D dynamic model based on a common central hotspot undergoing PdV work from a 3D shell consisting of individal 1D radial cold fuel profiles with separate areal densities and implosion velocities. The hotspot itself is 1D, with 0D power balance contributions from thermal conduction, alpha-heating and pure bremsstrahlung radiative loss. Using experimental inputs such as the implosion velocity, adiabat and 3D ρR variation from neutron diagnostics, the model was found to agree well with experimental observables such as yield and ion temperature.

Cheng *et al.* [86] quantified the effect of P_2 shape deformations on the hydrodynamic disassembly time of the hotspot by using the shortest radius of the deformed hotspot in the calculation. This was then used in a 1D ignition criterion [87] (explained further in section 2.4.1) to establish a shape-dependent ignition criterion. The analysis was also used in conjunction with other analytical models [87, 88] to compare the yield from these models to experimental data, to good agreement.

Taylor and Chittenden [89] instead considered this area from the viewpoint of how perturbations might affect the hotspot formation and conditions, using 3D simulations with local alpha-particle deposition and an escape-probability model for the radiation transport transitioning smoothly between optically-thin bremsstrahlung and optically-thick black-body emission. The study explored the volume-integrated hotspot power balance with an analytic model, and how perturbations affected the parameters of this model from their unperturbed values.

Bose *et al.* [90] also considered the impact of asymmetry on ICF hotspots in a similar manner. 2D *DEC2D* simulations were used without alpha-heating or radiation transport to explore the effect on hydrodynamics and thermal transport for sub-ignition implosion designs. The perturbations were found to reduce the *PdV* compression of the hotspot through residual mass flow. Hotspot bubbles were shown to be preferentially cooled by the ablative mass flow, reducing the hotspot burn volume in a manner reminiscent of the clean-volume analysis of Kishony and Shvarts [84].

Another study by Clark *et al.* [3] explored the requirements for achieving ignition using 2D and 3D HYDRA simulations. 3D simulations with a comprehensive set of perturbation sources (including fill-tube, fill-tube shadows, tent scar, radiation drive asymmetries and capsule surface roughness) were used to accurately recreate the experiment, and then 2D simulations were then hydrodynamically scaled to project the performance. It is noteworthy that these 3D simulations were the first HYDRA simulations of this type to include in-line Monte-Carlo alpha-particle transport, having previously been limited by computational memory constraints.

1.5 Overview of Thesis

In this work, we aim to explore the interactions of perturbations with the hotspot to investigate ignition and burn in their presence, considered from the viewpoint of the hotspot power balance. This work builds upon Taylor and Chittenden's work [89, 91] with upgraded physics models such as multi-group non-diffusive radiation transport and Monte-Carlo alpha-particle transport, and a focus on igniting and burning hotspots and the role of alpha-heating.

The background theory is given in chapter 2, in which we first describe the underlying physics of charged particle transport in high energy density plasmas; this is key to accurately modelling the energy deposition of alpha-particles in the hotspot and shell. The core physics of the hotspot power balance and perturbation growth via hydrodynamic instabilities are then given, along with a review of how perturbations can impact the hotspot formation and properties. The processes of ignition and burn in the hotspot are described, in addition to the theory behind hydrodynamic equivalence in scaling ICF implosions with capsule size and laser energy.

Accurate alpha-particle transport is necessary in order to ensure confidence in results within regimes where alpha-heating is the dominant source term. Numerical aspects of this are covered in chapter 3, in which we describe the Chimera code and the development of a Monte-Carlo charged particle transport model designed for the modelling of alpha-particles in ICF implosions. This includes computational method considerations as well as the implementation and testing.

One-dimensional modelling of the hotspot dynamics is explored in chapter 4, in which we examine the evolution of the hotspot with particular regard to the time- and radius-dependent variations of the power balance. The optical thickness of the hotspot is investigated, with Appendix A giving some additional background to radiative transfer

in plasmas. The properties of the hotspot and the power balance in various regimes of alpha-heating are compared. The effects of alpha-heating on the hydrodynamics of implosions scaled in capsule size and laser energy are shown in 1D, and provide insight into how the performance of the hotspot is dependent on the power balance. In addition, the dependence of these properties on how one defines the hotspot is studied.

Chapter 5 explores the impact of two idealised perturbations, a single spike and a single bubble, on the hotspot in 3D. The impact of these different types of perturbations on the various contributions to the hotspot power balance is examined and connected to their effects on the time-varying hotspot properties. The effects on overall performance are then discussed in relation to how ignition and burn are affected by the perturbations.

Chapter 6 investigates the scaling of 3D perturbed implosions with capsule size and laser energy. We examine how this scaling is affected by different forms of perturbations and by different implosion designs with different yield degradation levels. The effects of alpha-heating on the hydrodynamics under perturbed conditions are demonstrated, and the results of the scaling studies are compared to similar work projecting experimental performance as a function of driver energy.

Finally, chapter 7 summarises our findings and examines ways in which this work could be improved and built upon. In addition, potential avenues of research that are now available with the enhanced capabilities developed in the course of this work are suggested.

2 Theory

In order to understand hotspot ignition under inhomogeneous conditions, we must first consider the mechanisms through which the hotspot can gain or lose energy, and how these mechanisms are affected by asymmetry. Section 2.1 examines various models of charged particle transport in plasmas, in order to understand the process by which alpha-particles heat the hotspot. Section 2.2 then discusses the other mechanisms in the hotspot power balance, such as thermal conduction, radiation and mechanical work. The presence of perturbations can affect the hotspot power balance, and so section 2.3 considers the theory of perturbation growth and interactions with the hotspot. Following this, we then explore ways in which hotspot ignition is characterised in section 2.4, before explaining the various regimes the hotspot undergoes in its evolution towards propagating burn. Lastly, in section 2.5 the theory of hydrodynamic scaling of ICF capsules, which is used to access the various burn regimes, is given.

2.1 Charged Particle Transport

In this section we describe some commonly used models for the stopping of charged particles in plasmas, noting that an in-depth discussion of charged-particle stopping is beyond the scope of this work. It is instructive to briefly review a derivation of the classical Spitzer treatment, in which Coulomb collisions are considered with only small-angle scattering. We will then consider more complicated models which incorporate physics such as large-angle scattering, collective effects, quantum scattering, degeneracy and plasma coupling.

2.1.1 Spitzer

Following Spitzer [92], consider a particle of species α of velocity $\mathbf{v}_\alpha$, mass m_α and charge $Z_\alpha e$, colliding with background particles β of mass m_β and charge $Z_\beta e$. The long range of electrostatic forces results in the majority of collisions occurring at large distances with small-angle deflections, and thus we examine the effect of the sum of these small-angle collisions.

After N collisions, the change in velocity perpendicular to $\mathbf{v}_\alpha$ is given by the sum of the deflections:

$$\Delta v_{\alpha,\perp} = \sum_j^N (\Delta v_{\alpha,\perp})_j \tag{2.1}$$

The randomness of collisions means the exact value of $\Delta v_{\alpha,\perp}$ cannot be calculated, but

instead can be averaged for a large number of particles. This diffusive process results in an average $\langle(\Delta v_{\alpha,\perp})\rangle$ of zero for an isotropic distribution of velocities in field particles β, but a non-zero averaged diffusion coefficent of $\langle(\Delta v_{\alpha,\perp})^2\rangle$. For a Maxwellian, the non-zero coefficients are $\langle(\Delta v_{\alpha,\parallel})\rangle, \langle(\Delta v_{\alpha,\parallel})^2\rangle$ and $\langle(\Delta v_{\alpha,\perp})^2\rangle$.

For the case $m_\alpha \ll m_\beta$, we can calculate $\langle(\Delta v_{\alpha,\perp})^2\rangle$ from:

$$(\Delta v_{\alpha,\perp})^2 = \frac{4v_\alpha^2 b_0}{b} \tag{2.2}$$

where b is the impact parameter between two colliding particles, and $b_0 \approx e^2/4\pi\epsilon_0 m\langle v^2\rangle \approx e^2/(4\pi\epsilon_0 \cdot 3k_b T)$ is the impact parameter for a $90°$ deflection [93]. Averaging over test particles within a cylindrical shell defined by v_α and the impact parameter range $(b, b+db)$ gives:

$$\langle(\Delta v_{\alpha,\perp})^2\rangle = \int_{b_0}^{b_{max}} 2\pi b v_\alpha n_\beta (\Delta v_{\alpha,\perp})^2 db = 8\pi n_\beta v_\alpha^3 b_0^2 \ln\left(\frac{b_{max}}{b_0}\right) \tag{2.3}$$

where n_β is the number density of field particles, and the Coulomb logarithm $\ln\Lambda$ is $\ln(\frac{b_{max}}{b_0})$. Here, b_{max} is given by the Debye length $\lambda_e = \sqrt{k_b T_e \epsilon_0/e^2 n}$.

When generalised to any mass ratio m_α/m_β, Chandrasekhar [94] and Spitzer [92] give:

$$\frac{\partial v_{\alpha\parallel}}{\partial t} = -A(1 + m_\alpha/m_\beta)\frac{G(w)}{v_\beta^2} \tag{2.4}$$

$$\frac{\partial v_{\alpha\parallel}^2}{\partial t} = AG(w)/v_\alpha \tag{2.5}$$

$$\frac{\partial v_{\alpha\perp}^2}{\partial t} = A\frac{\text{erf}(w) - G(w)}{v_\alpha} \tag{2.6}$$

$$A = \frac{Z_\alpha^2 Z_\beta^2 e^4 n_\beta \ln\Lambda_{\alpha\beta}}{2\pi m_\alpha^2 \epsilon_0^2} \tag{2.7}$$

$$G(w) = \frac{\text{erf}(w) - w\frac{\partial\,\text{erf}(w)}{\partial w}}{2w^2} \tag{2.8}$$

where $w = v_\alpha/v_\beta$, $\ln\Lambda_{\alpha\beta}$ is the Coulomb logarithm between the two species, $\text{erf}(w) = 2/\sqrt{\pi} \times \int_0^w e^{-t^2} dt$ is the error function and $G(w)$ is the Chandresekhar function. We can calculate v_β as $v_\beta^2 = 2kT_\beta/m_\beta$. $\partial v_{\alpha\parallel}/\partial t$ is the frictional drag term stemming from the above coefficient $\langle(\Delta v_{\alpha,\parallel})\rangle$, while $\partial v_{\alpha\parallel}^2/\partial t$ and $\partial v_{\alpha\perp}^2/\partial t$ are the parallel and perpendicular diffusive terms respectively.

If we consider only the frictional term $\partial v_{\alpha\parallel}/\partial t$, the stopping power due to this term can be written simply as $dE/dx = m_\alpha dv/dt$. The full stopping power which includes energy-loss due to the diffusive terms is given:

$$\frac{dE}{dx} = \frac{Z_\alpha^2 Z_\beta^2 e^4}{4\pi\epsilon_0^2}\frac{n_\beta \ln\Lambda_{\alpha\beta}}{m_\alpha v_\alpha^2}\left[\frac{m_\alpha}{m_\beta}\psi(y) - \psi'(y)\right] \tag{2.9}$$

where $y = w^2$,

$$\psi(y) = \frac{2}{\sqrt{\pi}} \int_0^y \sqrt{\xi} e^{-\xi} d\xi = 2w^2 G(w) \tag{2.10}$$

and

$$\psi' = \frac{d\psi}{dy} = \frac{2}{\sqrt{\pi}} \sqrt{y} e^{-y} = \frac{2}{\sqrt{\pi}} w e^{-w^2} = w \frac{\partial \,\mathrm{erf}(w)}{\partial w} \tag{2.11}$$

2.1.2 Li and Petrasso

Li and Petrasso noted that the large-angle scattering events were likely important for stopping in inertial confinement plasmas, particularly at the end of the stopping range of an alpha-particle [12], and proceeded to generalise the Fokker-Planck equation to include the effects of large-angle scattering [95]. The impact of collective plasma oscillations is also included through an *ad hoc* "collective" Coulomb logarithm $\ln \Lambda_c$ [96] in their formulation:

$$\frac{dE}{dx} = \frac{Z_\alpha^2 Z_\beta^2 e^4}{4\pi\epsilon_0^2} \frac{n_\beta}{m_\beta v_\alpha^2} \left\{ H(y) \ln \Lambda_b + \Theta(y) \ln \Lambda_c \right\} \tag{2.12}$$

$$H(y) = \psi(y) - \frac{m_\beta}{m_\alpha} \left\{ \psi' - \frac{1}{\ln \Lambda_b} \left[\psi(y) + \psi' \right] \right\} \tag{2.13}$$

where $\ln \Lambda_c \equiv \ln \left[1.123 \, (y)^{1/2} \right]$, and $\Theta(y) = 0, 1$ for $y \leq 1, y > 1$ is the Heaviside step-function. Here, the collisional Coulomb logarithm is calculated as $\ln \Lambda_b = \ln(\lambda_D / b_{min})$, with a maximum impact parameter as the Debye length λ_D and a minimum $b_{min} = \sqrt{b_0^2 + \lambda_B^2}$, a quadrature sum of b_0, the impact parameter at $90°$ (with reduced mass m_r and relative velocity u) and the de Broglie wavelength $\lambda_B = \hbar/2m_r u$. In low temperature and high density regimes, λ_D and b_{min} calculations should include the effect of electron degeneracy. The large-angle scattering contribution is given by the $1/\ln \Lambda_b$ term in equation 2.13, while the collective effects are included by the $\ln \Lambda_c$ term in equation 2.12. Accordingly, equation 2.12 reduces to equation 2.9 without these corrections. A recent erratum [97] gives a correction to the model for v_β, which should be calculated using $k_b T / m_\beta$ rather than $2k_b T / m_\beta$ for the collective stopping power.

2.1.3 Maynard-Deutsch

The Maynard-Deutsch (MD) model [98] uses the random-phase approximation (RPA) [99] to treat the dielectric function in a plasma with arbitrary electron degeneracy. This model considers contributions from both free and bound electrons through Coulomb collisions and collective oscillations (plasmon modes), and allows slowing calculations for any velocity ratio v_α/v_e, but neglects ion contributions. The formulation uses Fermi-Dirac statistics to treat the degeneracy, but recovers the same results as the classical Boltzmann case at high plasma temperature. The overall effect of degeneracy is to reduce the stopping power, as compared to the non-degenerate case.

We will proceed to use Zimmerman's parameterisation of the model [100] (hereafter labelled as MDZ), which is significantly more computationally tractable as it does not

require numerical integration:

$$\frac{\partial E}{\partial x}\bigg|_e = \frac{Z_\alpha^2 e^4}{4\pi\epsilon_0^2}\frac{1}{m_e v_\alpha^2}n_F L_F \tag{2.14}$$

$$L_F = \frac{1}{2}\ln\left(1+\Lambda_F^2\right)\left(\text{erf}\left(w\right) - \frac{2}{\sqrt{\pi}}we^{-w^2}\right) \tag{2.15}$$

$$\Lambda_F = \frac{4\pi m_e v_e^2}{h\omega_{pe}}\cdot\frac{0.321+0.259w^2+0.0707w^4+0.05w^6}{1+0.130w^2+0.05w^4} \tag{2.16}$$

$$v_e = \begin{cases}\sqrt{\pi}h(2\pi m_e)^{-1}\left[4n_e\left(1+e^{-\mu/T_e}\right)^{1/3}\right] & \text{degenerate}\\ \sqrt{(2kT_e)/m_e} & \text{non-degenerate}\end{cases} \tag{2.17}$$

Zimmerman also includes a simple ion contribution, included here with a modification to the Coulomb logarithm [101]:

$$\frac{\partial E}{\partial x}\bigg|_i = \frac{Z_\alpha^2 Z_\beta^2 e^4}{4\pi\epsilon_0^2}\frac{n_\beta}{m_\beta v_\alpha^2}\ln\Lambda \tag{2.18}$$

$$\ln\Lambda = \frac{1}{2}\ln\left(1+\frac{b_{max}^2}{b_{min}^2}\right) \tag{2.19}$$

where the maximum impact parameter $b_{max} = \sqrt{\lambda_D^2 + a_i^2}$ for Debye length λ_D and ion sphere radius $a_i = (3/4\pi n_i)^{1/3}$, and the minimum impact parameter $b_{min} = \sqrt{b_0^2 + \lambda^2}$ for 90° impact parameter b_0 and de Broglie wavelength $\lambda = \hbar/2m_r u$ as defined previously.

2.1.4 Brown, Preston and Singleton

Brown, Preston and Singleton [102] use a dimensional continuation method to determine the charged-particle stopping for a weakly coupled plasma to leading order in plasma coupling parameter g:

$$g = \frac{Z_\alpha Z_\beta e^2 \kappa_\beta}{4\pi\epsilon_0 k_b T_\beta} \tag{2.20}$$

for a plasma species β of charge e_β, temperature T_β and inverse Debye length $\kappa_\beta = 1/\lambda_D$. This gives both classical and quantum contributions for dE/dx of the form $g^2\ln\left[Cg^2\right]$ for constant C, which is exactly calculable in their method, with no restrictions on particle charge, mass or speed (beyond being non-relativistic).

The general BPS stopping power consists of three contributions from the plasma background (for each species):

$$\frac{dE}{dx}_{BPS} = \left(\frac{dE_S^C}{dx} + \frac{dE_R^C}{dx}\right) + \frac{dE^Q}{dx} \tag{2.21}$$

where the first term corresponds to the classical short-range, Coulomb scattering, the second term is due to long-distance, collective effects and the last term represents short-range two-body quantum scattering.

The full expressions are not included here, but can be found in the appendix A in Singleton [103] or alternatively, equations (3.3), (3.4) and (3.19) in the Brown-Preston-Singleton paper [102].

The standard Spitzer derivation combines the short-range Coulomb scattering with the long-distance dielectric effects ($\mathbf{j} \cdot \mathbf{E}$ loss) through the impact parameter integral between the short-range minimum impact parameter b_0 up to the long-range maximum impact parameter of the Debye length (λ_D), producing the standard Coulomb logarithm factor $\ln \Lambda$, and some constants of order one. Brown, Preston and Singleton note that the computation of these constants that accompany this Coulomb logarithm factor is more difficult than reaching the logarithmic term itself, which can be used using relatively simple physical arguments. Their approach instead calculates Coulomb scattering for dimensions $\nu > 3$ and the $\mathbf{j} \cdot \mathbf{E}$ collective loss term for dimensions $\nu < 3$, thus avoiding any divergence in the respective integrals. The results are then obtained in the limit of $\nu \to 3$.

Quantum corrections are included, calculated by considering short-range quantum scattering effects. The importance of these effects can be gauged through the dimensionless Coulomb parameter $\eta_{\alpha\beta}$, which measures the strength of the interaction between the projectile particle α and the background particles β:

$$\eta_{\alpha\beta} = \frac{Z_\alpha Z_\beta e^2}{4\pi\epsilon_0 \hbar v_{\alpha\beta}} \tag{2.22}$$

where $v_{\alpha\beta} = |\mathbf{v}_\alpha - \mathbf{v}_\beta|$ is the (average) relative velocity. When $\eta_{\alpha\beta} \ll 1$ (i.e. $v_{\alpha\beta} \gg 1$), quantum scattering is significant.

2.1.5 T-Matrix

Gericke *et al.* [104–106] considered stopping power using the quantum Boltzmann equation for two-body interactions, under both static [104] and dynamically-screened [106] Coulomb potentials. The dynamic screening adjusts the Coulomb potential based on a velocity-dependent effective screening length, $\lambda(v)$ with a term $\exp(-r/\lambda(v))$. The model name stems from the T-matrix in the Lippmann-Schwinger equation describing the two-particle quantum scattering interaction. The approach thus considers binary collisions of arbitrary strength (i.e. both small-angle and large-angle scattering). While the static case requires collective effects to be added through a first-order RPA solution to the Lenard-Balescu equation, the dynamic-screening includes these effects *ab initio*. However, the approach is not rigorous in that the screening length $\lambda(v)$ is adjusted such that the overall stopping power tends towards the correct asymptotic limits.

2.1.6 Theoretical Comparison

In comparing the validity and accuracy of various stopping models, we need to consider both the plasma conditions for which the model is valid, as well as the physical effects that the model seeks to address. The plasma conditions under which the stopping power

is calculated can be measured by the degeneracy parameter, $\Theta = k_b T/E_F$, for Fermi energy E_F and the plasma coupling, either through g (see equation 2.20) or Γ, where $g^2 = 3\Gamma^3$; g uses the Debye length whereas Γ uses the Wigner-Seitz radius. The strength of the coupling between the charged particle and the plasma can be measured by the quantum parameter $\eta_{\alpha\beta}$, as defined above in equation 2.22. The physical interactions between the particle and the background plasma included in the stopping model can include short-range small-angle and large-angle Coulomb scattering, long-range collective oscillations, and quantum scattering.

Clearly, the Spitzer model is the simplest scenario, considering only classical small-angle scattering in an ideal plasma. The Chandrasekhar function varies according to v_α/v_β, and therefore accounts for some variation in the particle-plasma coupling $\eta_{\alpha\beta}$. We note that effects such as degeneracy or quantum scattering can be included *ad hoc* through the Coulomb logarithm, such as by using a quadrature sum of the electron temperature, T_e, and the Fermi temperature, T_F, in the calculation of the Debye length, or a quadrature sum of the de Broglie length and b_0 as the minimum impact parameter. LP includes both small-angle and large-angle Coulomb scattering, as well as an *ad hoc* collective oscillation term. As mentioned for Spitzer, the effects of degeneracy and quantum scattering can be included in an *ad hoc* manner via the Coulomb logarithm and b_{min} calculations respectively. In addition, the variation in $\eta_{\alpha\beta}$ is accounted for in a similar manner as in Spitzer, using the Chandrasekhar function in a modified form. MD includes small-angle Coulomb scattering and long-range collective oscillations, at arbitrary degeneracy. The result includes the effect of quantum scattering through the quantum dielectric function, but not large-angle scattering. The MD calculation is derived for arbitrary particle velocity, which is parameterised in a thermal coupling factor similar to the Chandrasekhar function, like in the Spitzer and LP models. However, it does not consider large-angle scattering or strong plasma coupling. BPS includes small-angle Coulomb scattering, collective oscillations and quantum scattering, to all orders of $\eta_{\alpha\beta}$ for weakly-coupled plasmas, but does not include large-angle scattering or consider degeneracy. The T-matrix approach considers Coulomb scattering beyond first-order, including large-angle scattering and strong coupling, but uses a simple, *ad hoc* dynamic screening length to incorporate the effects of collective oscillations and does not include degeneracy.

2.1.7 Experimental Comparison

The Maynard-Deutsch and BPS have shown very similar results in the weakly coupled ($\Gamma \leq 0.1$) plasma regime [107]. In the less weakly coupled regime ($0.1 < \Gamma \lesssim 1$), Hayes *et al.* [108] experimental study using reaction-in-flight (RIF) spectra measured on NIF implosions found Maynard-Deutsch to agree the best with the measurements. Since RIFs occur mainly in the cold fuel, the study focuses predominantly on constraining stopping powers under these strongly degenerate and coupled conditions ($n \sim 10^{32}m^{-3}, T \sim 0.5keV \approx \Gamma \sim 0.2, \Theta \sim 0.6$). The BPS model was not used as the plasma conditions

were beyond the validity of the model. The LP model had significantly worse agreement with the experimental results here. However, it is unclear that these LP calculations included the correction [97], which should have brought the model into better agreement with MD.

Frenje *et al.* [109] explored the charged-particle spectra using an OMEGA implosion of equimolar D-He^3, with plasma conditions ranging from $n_e \sim 10^{28} - 10^{29} m^{-3}$ and $T_i \sim 3 - 12 keV$ ($\Gamma \ll 1$ and $\Theta \gg 1$). Spectra were measured for DD tritons, DD protons, DHe^3 alphas and DHe^3 protons, and found that the BPS model required the quantum corrections to agree with experimental data. The data were found to generally agree well with both the BPS and LP (including the correction) models, particularly for $v_i > v_{th}$; here, v_i is the velocity of the charged particle and v_{th} the thermal plasma electron velocity However, due to the lack of a T_e measurement, the study was unable to definitively discriminate between the BPS and LP models, which demonstrated differences of $\sim 20\%$ around $v_i \sim v_{th}$.

Cayzac *et al.* [110] also compared various stopping models with experimental data in fully ionized plasmas at $n_e \approx 5 \times 10^{26} m^{-3}$ and $T_e \approx 150 eV$, corresponding to weak plasma coupling $\Gamma \approx 0.01$ and low degeneracy. The study focuses on the energy loss around the Bragg peak, i.e. $v_i \sim v_{th}$, and shows good agreement of BPS and a T-matrix approach [106] with the experimental data. The LP model (without the correction) is also compared to the T-matrix model, and appears to overpredict the stopping power by around $20 - 25\%$.

Note that the inclusion of quantum effects on the stopping power, such as quantum scattering and degeneracy, is significant, and can result in a 30% reduction in the overall stopping power [107–109].

It seems that the BPS and the T-matrix models agree quite well with available experimental data within the range of validity (weak coupling). The corrected LP model also appears to give better agreement than the uncorrected version, but it is still unclear as to whether agrees as well as BPS. Our model (described later in chapter 3) uses the MDZ model as described above for the stopping, in combination with the Spitzer diffusion coefficients. The MDZ model is simpler and faster to compute than either the BPS or the T-matrix model, and it will be shown later (see section 3.5) that the MDZ model agrees well with the BPS model, which agrees with experimental data.

2.2 Hotspot Power Balance

The power balance within a hotspot determines whether a hotspot is self-heating, which is necessary for ignition to occur, and is dictated by the equation:

$$\frac{dE}{dt} = W_{dep} + W_m + W_e + W_\gamma \tag{2.23}$$

where W_{dep} is the power density contribution due to deposition of fusion products, W_m due to mechanical work, W_e due to thermal conduction and W_γ due to radiation. Each
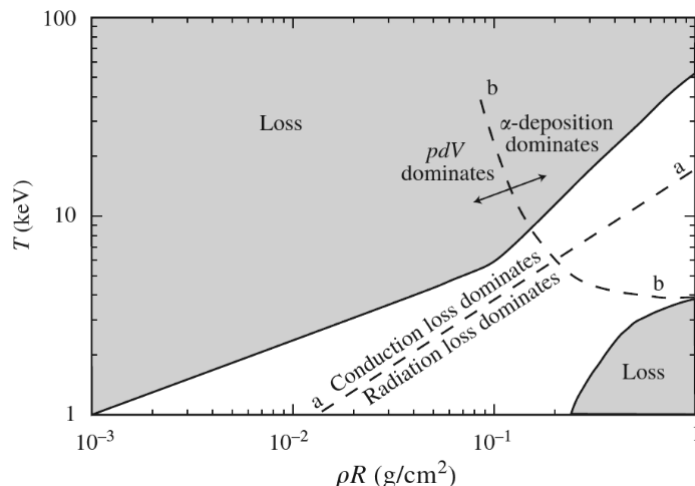
Figure 2.1: Hotspot energy loss and gain regimes in $\rho R - T$ space, shown for an implosion velocity of $u = 300 kms^{-1}$. Reprinted from [6] with the permission of AIP Publishing.

of these terms will be explained in more detail in the following subsections.

Figure 2.1 shows the regions in areal density-temperature parameter space in which each of the processes dominates the hotspot power balance, for an example scenario where the implosion velocity is equal to $300 kms^{-1}$. The two loss regions are radiation and conduction loss dominated, while the gain region can be divided into two: the hotspot creation region dominated by mechanical work conversion to internal energy, and the self-heating region, dominated by deposition of fusion products (normally alpha-particles). As the implosion progresses, the hotspot follows a trajectory moving from low ρR and T to higher values of both through different regimes – and ideally into the self-heating region. The boundaries as shown in figure 2.1 [6] are an example, and will vary as the implosion velocity changes throughout the implosion.

2.2.1 Mechanical Work, W_m

Mechanical work is responsible for the creation of the hotspot, with the imploding shell acting as a piston to compress and heat the inner region. In general, the change in energy due to work being done by a volume at a pressure, P, with a resultant change in volume, dV, is $dE = PdV$. Hence for a hotspot being formed by an imploding shell, the volume-integrated mechanical work on the hotspot is given:

$$W_m = \frac{dE}{dt} = -P\frac{dV}{dt} = PSv \tag{2.24}$$

for a sphere of radius r, volume V and surface area S, imploding at velocity v. For a hotspot of radius $R_{hs}(t)$, v is equal to the fluid velocity at $R_{hs}(t)$, not $dR_{hs}(t)/dt$. This term contributes towards heating the hotspot until stagnation, at which point the shell can no longer do work against the hotspot pressure. Once the hotspot starts to expand

against the inertia of the shell, it loses energy as work is instead done on the shell.

2.2.2 Thermal Conduction, W_e

Thermal conduction losses from the hotspot are predominantly due to electron conduction — ion conduction is much less efficient at transporting heat due to the larger ion mass. Following [89], the thermal conduction losses from the hotspot can be calculated in 1D from the divergence theorem:

$$W_e = -\int \nabla \cdot (\kappa \nabla T) dV = -\int \kappa(R_{hs}) \left.\frac{dT}{dr}\right|_{r=R_{hs}} dS \tag{2.25}$$

for a hotspot of radius R_{hs} with surface and volume elements dS, dV respectively. Assuming a classical Spitzer conductivity for electron-ion collisions [92]:

$$\kappa_{Spitz} = \left(\frac{8}{\pi}\right)^{\frac{3}{2}} \left(\frac{1}{1 + \frac{3.3}{Z}}\right) \frac{(k_B T)^{\frac{5}{2}} k_B}{Z e^4 \sqrt{m_e} \ln \Lambda} \tag{2.26}$$

where $\ln \Lambda$ is the electron-ion Coloumb logarithm.

For a 1D implosion, the temperature profiles of the hotspot and cold shell can be described by [89]:

$$T(r,t) = T_{hs}(t) - (T_{hs}(t) - T_R(t))\left(\frac{r}{R_{hs}(t)}\right)^q \tag{2.27}$$

where the exponent q can be species-dependent, T_{hs} is the temperature at the centre of the hotspot and T_R is the temperature at the edge of the hotspot (i.e. at $r = R_{hs}$). Using equation 2.27 for dT/dr, replacing the constants in equation 2.26 with a single constant A and integrating for the spherically symmetric, 1D case:

$$W_e = \frac{A_{TC} R_{hs} q (T_{hs} - T_R) T_R^{\frac{5}{2}}}{\ln\Lambda} \tag{2.28}$$

Although the thermal conduction transports heat out of the hotspot into the shell, this process actually both drives mass ablation of shell material back into the hotspot, and heats up this ablated mass. In this manner, less energy is required from the hotspot itself to heat up the ablated mass and therefore heat flow from the hotspot can be seen as being recycled into the hotspot as mass flow, rather than purely as energy lost [111].

2.2.3 Radiation, W_γ

For low-Z plasmas at high temperatures of several keV, the primary radiative mechanism is bremsstrahlung. If optically-thin, the radiative losses would be equal to the frequency-integrated bremsstrahlung emissivity, which is given [112]:

$$J = \int_0^\infty j_\nu d\nu = \frac{32\pi}{3}\left(\frac{2\pi k_b T_e}{3m_e}\right)^{\frac{1}{2}} \frac{Z^2 e^6}{m_e c^3 h} n_i n_e \tag{2.29}$$

for a plasma with electron temperature, T_e, and electron and ion number densities n_e, n_i respectively. This means that the emissive power of a given optically thin volume can be given as:

$$W_\gamma = A_\gamma \int (\rho(r))^2 (T(r))^{1/2} dV \tag{2.30}$$

The density profile can be related to the temperature profile given in equation 2.27, assuming isobaricity. Ignoring any reabsorption, the integral can be approximated to give a profile factor of $\sim T_{hs}^{0.7}$ (the details of which can be found in the thesis of S. Taylor [91], equations $(3.3.9) - (3.3.12)$), finally giving:

$$W_\gamma \approx A_\gamma \rho_{hs}^2 T_{hs}^{1.2} V \tag{2.31}$$

where T_{hs} is in keV.

This provides a useful indication of how the radiative losses within the hotspot scale. However, as will be shown in section 4.2, the hotspot is not entirely optically thin — $1 keV$ photons in hotspot conditions have a mean free path of $\sim 25 \mu m$, which is of the order of the size of the hotspot. This indicates that some of the radiation emitted from the hotspot will be reabsorbed as it transits across the hotspot. Indeed, considering the expression for the frequency-dependent bremsstrahlung absorption [112]:

$$\kappa_\nu = \frac{4}{3} \left(\frac{2\pi}{3m_e k_b T} \right)^{\frac{1}{2}} \frac{Z^2 e^6}{hc m_e \nu^3} n_i n_e \tag{2.32}$$

we can see that the cold dense shell will strongly absorb radiation from the hotspot, due to the $n_i n_e$ scaling.

2.2.4 Fusion Product Deposition, W_{dep}

In a DT plasma, the fusion reaction $D + T$ produces $3.45 MeV$ alpha-particles and $14.1 MeV$ neutrons. Of these products, the main interaction for the neutrons is elastic scattering with the plasma ions, losing (on average) a fraction $2A/(A+1)^2$ of their energy during a collision with an ion of mass number A. The mean free path of the neutrons is given by $l_n = 1/\sigma n$ for number density n and collision cross-section σ, which is equal to $0.8 \times 10^{-28} m^2$ for D and $10^{-28} m^2$ for T. This gives $l_n \approx 1.1 \times 10^{-3} m$, significantly larger than the hotspot radius $R_{hs} \approx 3 \times 10^{-5} m$. Therefore, we do not include neutron energy deposition in our studies of central hotspot ignition. It is worth noting that neutron energy deposition can have an effect on the evolution of the propagating burn wave, and any potential detonation transition [113].

In addition, the effect of the plasma ambipolar E-field, given by:

$$e n_e \mathbf{E} = -\nabla \mathrm{P} - \beta n_e \nabla T \tag{2.33}$$

can be ignored for alpha-particle transport in hotspots. Assuming an isobaric, unmagnetised DT fuel assembly (such that $\nabla \mathrm{P} = 0, \beta = 0.7029$ [114]), the electric field can be

calculated by:

$$e\mathbf{E} = -\beta\nabla T \tag{2.34}$$

For a hotspot of size $\sim 30\mu m$ in radius, central temperature $7keV$ and edge temperature $1keV$, the energy gained by an alpha-particle in this electric field $\sim 8.5keV$. Since this is much less than the $3.45MeV$ of the alpha-particle, the effect of this is negligible for the dynamics of alpha-heating.

The alpha-particles deposit energy back into the plasma via Coulomb collisions, with differing fractions of their total energy being coupled back into the electrons and ions. The stopping power of an alpha-particle on a plasma species depends on the ratio of the velocities between them — the closer the ratio to 1 (i.e. the closer the velocities), the stronger the interaction. Hence, the fraction of energy coupled into each species is highly temperature-dependent — at lower temperatures the majority of the Coulomb collisions occur with the electrons, whereas as T_e increases, the scattering due to (and therefore contribution from) ions becomes more significant. To illustrate, the velocity of a $3.45MeV$ alpha-particle is $\sim 10^7 m/s$, $\sim 3 \times 10^7 m/s$ for a $5keV$ electron, and $\sim 5 \times 10^5 m/s$ for a $5keV$ DT ion. At solid density ($\rho = 225kgm^{-3}$), the fraction deposited into the ions is approximately [12]:

$$f_i \approx \frac{1}{1 + \frac{32}{T_e}} \tag{2.35}$$

This gives the temperature at which equal fractions are deposited into each of the electrons and ions as $T_e = 32keV$. Note that this temperature has a small density dependence, and decreases to $T_e = 24keV$ at $\rho = 10^7 kg/m^3$ [12]. Since the temperatures here are significantly lower than 25-30keV, the majority of the energy will be deposited into the electrons, which then equilibrate with the ions. The alpha-particle velocities decrease as [12]:

$$\frac{dv_\alpha}{dt} = -\frac{v_\alpha}{2t_{\alpha e}}, \qquad t_{\alpha e} \simeq \frac{42T_e^{\frac{3}{2}}}{\rho\ln\Lambda_{\alpha e}}ps \tag{2.36}$$

for electron temperature T_e in keV, density ρ in g/cm^3, and Coulomb logarithm $\ln\Lambda_{\alpha e}$ between alpha-particles and electrons. Considering the eqns. 2.36, it is clear that the cold, dense shell outside the hotspot is very effective at stopping the alpha-particles, and thus there is a thin layer in the hotspot-shell boundary which is rapidly heated. This is the process by which burn waves begin to propagate post-ignition.

2.3 Perturbations

Hydrodynamic instabilities, of which the Rayleigh-Taylor Instability (RTI) is particularly prevalent, can greatly exacerbate the small imperfections inherently present within the capsule over the course of the implosion.. In this section, we describe the growth rate of the RTI and its relevance to ICF, in addition to how perturbations can affect the hotspot.

2.3.1 Perturbation Growth

The RTI occurs at the interface between two fluids of differing densities (ρ_1, ρ_2, where $\rho_1 > \rho_2$), where the heavier fluid is being accelerated into the lighter fluid at some rate, g. For the simplest case [115], a perturbation to the interface would grow exponentially with time, $\sim e^{\gamma t}$, with the rate given by:

$$\gamma = \sqrt{A_T k g} \tag{2.37}$$

for Atwood number $A_T = \sqrt{(\rho_1 - \rho_2)/(\rho_1 + \rho_2)}$ and perturbation wavenumber k.

2.3.1.1 Ablative Rayleigh-Taylor Instabilities

The ablative RTI is a modifed version of the above scenario, in which the interface is undergoing some kind of ablation, and therefore can be seen as a moving interface. The laser-driven ablation of the outer surface of an ICF capsule undergoes precisely this situation, and therefore is susceptible to the ablative Rayleigh-Taylor instability.

It has been shown [80, 116–119] that ablation at the interface reduces the instability growth rate, the effect of which can be included in an extension to the growth rate in equation 2.37 [117] to give:

$$\gamma = \alpha\sqrt{A_T k g} - \beta k v_a \tag{2.38}$$

for ablation velocity v_a and constants α, β, which are of order 1. Instinctively this can be seen as the removal of material from the edge of the unstable perturbation due to ablation, which therefore reduces the size of the perturbation.

This can be further extended to an interface of finite scale length, $L_m = min\,|\rho/\left(d\rho/dx\right)|$, within which the density varies smoothly between the two values. The growth rate is then given by an extension to equations 2.37 and 2.38 [120]:

$$\gamma = \alpha\sqrt{\frac{A_T k g}{1 + A_T k L_m}} - \beta k v_a \tag{2.39}$$

Physically, the finite scale length reduces the density gradient and therefore the density discontinuity characterised by the Atwood number A_T.

2.3.1.2 Deceleration Phase Rayleigh-Taylor Instabilities

The ablative RTI is most commonly considered during the ablation phase of the implosion, where the ablator material is ablated by radiation and accelerates the capsule inwards. However, it is also relevant on the inner surface of the shell during the deceleration phase, during which the dense outer shell is being decelerated by the rapidly increasing hotspot pressure. The deceleration phase RTI features significant ablative stabilisation [80, 111] due to heat flux from thermal conduction, radiation transport and alpha-transport out of the hotspot. Sources of perturbations here can arise due to surface defects on the inner surface, in addition to those which "feed-through" from the

outer surface of the shell (due to, for example, radiation asymmetries or the tent scar).

2.3.2 Impact on the hotspot

While the heat flow from the hotspot can affect the size and growth of perturbations through ablative stabilisation, the perturbations themselves also affect the energetics and evolution of the hotspot.

Perturbations can affect the implosion velocity and synchronisation of fuel stagnation, and therefore the overall efficiency of conversion from kinetic energy to thermal energy [89], resulting in residual kinetic energy (RKE) in the fuel due to unstagnated mass flow within the hotspot [40]. This can be illustrated through pressure differences between the hotspot and perturbed regions. The spike regions take longer to decelerate and penetrate deeper into the hotspot owing to the high ram-pressure compared to the hotspot thermal pressure, while the lower ram-pressure bubble regions stagnate and re-expand earlier against the hotspot thermal pressure. As a result, the KE conversion to thermal energy is spread over a longer time period owing to the varying velocities of different regions. If the spike amplitude is high, the KE contained in the spike may not all be thermalised, unless a collision occurs with another spike in the centre of the hotspot. In the unperturbed case, the entirety of the shell stagnates simultaneously, and therefore the conversion of KE to thermal energy occurs in a shorter, higher peaked impulse.

In general, the perturbations can increase the thermal conduction losses from the hotspot due to both temperature gradient compression and increasing overall surface area of the hotspot. If large-enough amplitude, the perturbation could entrain high-Z ablator and dopant material into the hotspot (as the fill-tube does [8]), which would increase radiative losses. The average implosion velocity of a perturbed shell can increase owing to mass flow from spikes into bubbles [89], and in combination with a reduction in temperature from the less efficient PdV heating, can result in higher hotspot densities and therefore radiation emission.

Taylor and Chittenden [89] characterised perturbations into 'shape' and 'mix' perturbations, as defined by the characteristic size of the hotspot bubble, R_b and the rate of thermal conduction collapse of the bubble, which scales as $\sim R_b^2$. For shorter-wavelength (smaller-scale) mix perturbations, hotspot bubble collapse occurs on a faster timescale than the stagnation time. The thermalisation of the bubble material with the shell (i.e. the mixing of the hotspot and shell material) removed thermal energy from the hotspot and reduce the total hotspot volume. Longer-wavelength shape perturbations result in bubbles which remain part of the hotspot and therefore hot compared to the shell. These bubbles increase overall thermal conduction losses but at a slower rate than for the mix perturbations, and does not reduce the hotspot volume. These can be considered as an effective reduction in the hotspot radius [84]. Different modes of perturbation can also interact with one-another, resulting in off-centre trajectories and therefore worse KE thermalisation [89].

Note that even with highly perturbed hotspots, early time perturbation growth does more work on the hotspot and therefore boosts the early time fusion yield even as the later-time fusion burn is quenched.

2.4 Hotspot Ignition and Burn

In this section, we describe the hotspot evolution towards ignition and burn after the formation at stagnation.

2.4.1 Ignition

In magnetic confinement fusion, ignition is defined as the point where fusion reaction energy-deposition compensates entirely for the losses, and the plasma requires no more energy input to sustain fusion reactions. This definition is usually formulated in the form of a Lawson-type criterion [121], such as $P\tau > f(T)$.

However, the non-steady-state, transient nature of ICF means that this MCF definition is not directly applicable. Levedahl and Lindl [83] defined ignition as "when energy deposition from the central hotspot is sufficient to generate a self-sustaining burn wave in the surrounding fuel". In their following investigation of scaling laws for ignition, they defined the point of 'marginal ignition' as a capsule gain of 1 — defined as the ratio of the capsule yield to the (estimated) energy absorbed by the capsule. Scaling laws for this ignition energy, E_{ign} have been explored under the conditions of constant drive pressure, varying implosion velocity and fuel adiabat [83, 122], presented as $E_{ign} \sim \alpha_{if}^a v_{imp}^b$ for in-flight adiabat α_{if} and implosion velocity v_{imp}, and constants a, b. Herrmann *et al.* [123] explored a similar case for varying drive pressure P, resulting in an ignition condition of the form $E_{ign} \sim \alpha_{if}^a v_{imp}^b P^c$.

Betti *et al.* [124] instead modified a simple Lawson criterion derived from a steady-state balancing of alpha-particle heating with energy losses:

$$P\tau > \frac{24}{\epsilon_\alpha S(T_0)} \tag{2.40}$$

to define ignition as $\chi = 1$ for the ignition parameter:

$$\chi = \frac{1}{24} P\tau \epsilon_\alpha S(T_0) \tag{2.41}$$

where

$$S(T_0) \equiv \frac{1}{V_{hs}} \int_{V_{hs}} \frac{\langle \sigma\nu \rangle}{T^2} dV \tag{2.42}$$

and $\epsilon_\alpha = 3.5 MeV$ is the alpha-particle birth energy. Here, the confinement time τ_{conf} is defined as R_s/v_{imp}, for cold shell radius R_s and implosion velocity v_{imp} (mass-averaged across thickness of the shell). The ignition parameter definition is manipulated into a form based on experimentally measurable parameters: the neutron averaged total areal density $(\rho R)_{tot}$ and temperature T. The condition can be extended into 3D using

a 'clean' volume analysis [84], where the hotspot volume is effectively reduced due to perturbations.

Cheng *et al.* [87, 125] derive an alternative ignition criterion by requiring the fusion reproduction time, τ_{rep}, to be less than the confinement time, τ_{conf}, taken to be hydrodynamic disassembly time of the hotspot, τ_H ($= \tau_{conf}$). This definition of the confinement time is one of the key differences to Betti's model; Cheng *et al.* define it as the hydrodynamic disassembly time of *the hotspot*, rather than the entire capsule:

$$\tau_H = \frac{R_{hs}}{C_s^*} \tag{2.43}$$

for hotspot radius R_{hs} and hotspot effective sound speed, $C_s^* = C_s/f_T$; this is the hotspot adiabatic sound speed, C_s, adjusted by the cold fuel tamping through tamping factor $f_T \propto \sqrt{\rho_{shell}/\rho_{hs}}$ [7]. The fusion reproduction time is calculated as:

$$\tau_{rep} = \frac{E_T}{\dot{E}} \tag{2.44}$$

where the total hotspot energy density is:

$$E_T = 3(n_D + n_T)k_b T + E_{rad} \tag{2.45}$$

E_{rad} is the radiation energy density, and the hotspot power balance is:

$$\dot{E} = n_T n_D \langle \sigma\nu \rangle_{DT} W_\alpha - \sum_i \dot{Q}_l^i \tag{2.46}$$

which includes fusion product energy deposition (the first term) as well as energy loss terms, Q_l^i. However, for simplicity, the electron thermal conduction losses and bremsstrahlung were ignored, and a simple black-body radiation loss used ($E_{rad} = aT^4$) to give an ignition criterion in the form:

$$(\rho R)_{hs} \geq \frac{(1+d)^2}{d} \frac{\left[3k_b T + aT^4/n_{DT} \right] C_s A_{DT}}{f_T \langle \sigma\nu \rangle_{DT} W_\alpha N_A} \tag{2.47}$$

where $d = n_D/n_T$ is the deuterium:tritium ratio, $n_{DT} \equiv n_D + n_T = \rho_{DT} N_A/A_{DT}$, A_{DT} is the atomic mass of the DT mixture, and $\rho_{DT} = \rho_{hs}$ here, as n_D, n_T refer to the number densities within the hotspot. This shows the other key distinction of the model, which is the use of the hotspot areal density rather than the total areal density. The physical reasoning here is that a necessary and sufficient condition for ignition to occur is a sustained thermonuclear burn in the hotspot, and therefore the areal density of interest is that of the hotspot rather than the total areal density. These differences result in a stricter ignition criterion.

Springer *et al.* [85] recently proposed an ignition criterion based on incorporating the effects of PdV expansion on the cooling of the hotspot, not just the positive PdV work during compression. The key criteria for ignition was identified as: the net alpha-

heating in the hotspot must be both positive and increasing faster than the rate of PdV expansion. This criterion is equivalent to requiring $\frac{dT}{dt} > 0$ and $\frac{d^2T}{dt^2} > 0$ at minimum volume.

2.4.2 Burn Regimes

More qualitatively, we can examine the different characteristics of the hotspot when it undergoes various phases/regimes of 'burn'. Here, we describe briefly the properties in each regime, but reserve a more complete illustration and discussion for chapter 4.

In the non-igniting or '**self-heating**' phase, the alpha-particles produced by relatively weak levels of fusion heat the capsule more, depositing their energy throughout the hotspot. Although the centre of the hotspot is the hottest and therefore produces the most alpha-particles, the non-local alpha-particles are emitted isotropically and therefore results in similar levels of heating throughout the hotspot. The hotspot temperature remains relatively low during this phase, and therefore the hotspot transparency to alpha-particles is relatively low, owing to higher stopping within the hotspot; here, we define the transparency as $\tau_\alpha = l_\alpha/R_{hs}$ for alpha-particle mean free path l_α and hotspot radius R_{hs} [126, 127]. During stagnation, this self-heating phase begins in tandem with the PdV work being done on the hotspot. Failure can occur at this phase if the timescale of the alpha-heating is too long relative to the confinement time — in other words, if the capsule takes longer to heat up from fusion reactions than is required for the capsule to disassemble hydrodynamically.

The temperature increase of the hotspot due to alpha-particle energy deposition results in a significant boost to the fusion reactivity. Within the temperature range $2 - 6keV$, the fusion reactivity scales roughly as $\sim T^4$. Thus the increase in reactivity in turn produces more fusion, which produces more alpha-particles; this is the positive feedback/bootstrapping process by which high gain in central ignition inertial confinement fusion occurs. We term this the '**robust ignition**' phase. This increase in temperature also produces a corresponding increase in pressure, which can lead to earlier capsule re-expansion and can lead to failure if inadequately confined. The stopping power within the hotspot decreases as the temperature increases, and thus the mean free path of the alpha-particles increases, depositing more energy into the hotspot-shell boundary layer and the shell itself.

Although there is a degree of alpha-particle energy deposition into this layer during the robust-ignition phase, it is the significant increase in this deposition (due to the hotspot-alpha transparency increasing) which begins the '**propagating burn**' phase. As the energy deposition into the shell increases, more material gets ablated off the shell. This begins the propagation of the deflagration wave, as the front of alpha-particle energy deposition moves into the cold fuel layer, ablating it away. Deflagration is a sub-sonic, ablation-driven burn wave, in which the heat due to electron thermal conduction and alpha-heating flows out of the hotspot into the shell, and mass flows from the shell into the hotspot. This process increases the total mass of burning material,

and thereby boosts the yield again. Note that the hotspot density does not necessarily increase during this burn-wave propagation, due to the increase in volume through both mechanical expansion and the propagation of the burn-wave. The build-up of hotspot pressure during this phase accelerates the re-expansion, and can exacerbate weak points of confinement into an aneurysm-like path to failure.

During deflagration, the deposition of alpha-heating and electron thermal conduction energy is correlated. However, the hotspot is becoming progressively hotter and producing more fusion alpha-particles, to which it is also becoming more transparent. This results in an intensification of the alpha-particle energy deposition into the cold fuel, accelerating the deflagration front and compressing the cold fuel shell even more. This compression can lead to a shockwave which transitions into a '**detonation**' wave [113]; a detonation wave is a supersonic burning shockwave in which the propagation of the shock through material raises the temperature enough to ignite it. Detonation can result in a large yield gain due to the elevation of the entirety of the fuel into a burning plasma state, in which it continues to fuse during the post-detonation disassembly.

Gauthier *et al.* [113] describe an alternative path to the deflagration-detonation transition, in which neutron energy deposition reduces the deflagration-wave compression, but allows the alpha-particles to penetrate deeper into the fuel. The pre-heating of the cold fuel by the neutrons allows the alpha-driven deposition front to propagate ahead of the electron thermal conduction front. The alpha-driven deposition front accelerates up the density gradient as the reaction time in the material decreases sharply owing to the rapid boost in temperature. The subsequent deceleration of the deflagration wave as it propagates down the density gradient (past the peak of the shell density) allows the plasma expansion to produce a shockwave. The amplification of this shockwave through the rest of the shell then transitions into a detonation front. Both of these deflagration-detonation transitions are effectively due to the acceleration of the deflagration front into a supersonic wave.

2.5 Hydrodynamic Scaling

The theory of hydrodynamic equivalence in the scaling of ICF hotspot ignition implosions was developed by Nora *et al.* [128] in order to compare implosions between the 30kJ OMEGA facility [14] (which can perform smaller implosions up to multiple times a day) and the 1.8MJ NIF facility (which performs larger implosions, but at most once a day).

A set of hydrodynamically-equivalent implosions can be defined as having the same: unablated mass fraction, in-flight aspect ratio (both initial and during the acceleration phase), dimensionless trajectory, and dimensionless thickness. Using a 1D rocket model to model the implosion, hydro-equivalence can be shown to require the same implosion velocity V_{imp}, adiabat α and ablation pressure or drive intensity, assuming the same kind of pulse shape.

An implosion with initial radius R_0 requires total laser energy $E_L \propto R_0^3$ to drive it,

with a laser power $P_L \propto R_0^2$, implosion timescale $t \propto R_0$. This means that if one were to scale up a particular capsule design by some scale factor S, then the driver energy would need to increase by a factor of S^3. In this scaling, all of the capsule radii and shell thicknesses are therefore increased by a factor S. In the case of indirect-drive, this would also require an increase in hohlraum size by S so as to maintain an identical case-to-capsule ratio (CCR).

The hydrodynamic scaling is also expected to apply to the acceleration-phase RTI, with the growth factors for all modes remaining identical due to the capsule evolving to the same convergence ratio from the same dimensionless radius. However, in order to obtain the correct ablation front properties and achieve the right implosion velocity, aspects of the capsule such as the dopant concentration and ablator thickness would have to be reduced [3]. This is because various properties of the drive do not scale hydrodynamically, such as the optical depth for x-ray absorption.

In addition, deceleration-phase RTIs also do not scale hydrodynamically, as the ablative heat flow from the hotspot also does not scale hydrodynamically. As seen in section 2.2, the ablative heat flow involves alpha-heating, thermal conduction and radiation transport. This can be seen first from scaling of the hotspot temperature, which is not expected to be invariant; in an analytic adiabatic model of the hotspot neglecting alpha-heating and radiative losses, and considering thermal conduction losses as completely recycled due to ablation, Zhou and Betti calculated the hotspot temperature to scale with the laser energy as $\sim E_L^{2/21}$. This was a reasonable match to the fit calculated from 1D LILAC simulations of $\sim E_L^{0.07}$ (i.e. $\sim S^{0.2}$) [129]. From the hotspot temperature scaling ($\sim S^{0.2}$), and using equations 2.28 and 2.30, it can be seen that the thermal ($T^{7/2}$) and radiative ($T^{1/2}$) losses from the hotspot are also not scale invariant, but scale roughly as $\sim S^{0.7}$ and $\sim S^{0.1}$. Although the radiative contribution to the ablative heat flow is a smaller proportion of this optically-thin volume emissivity calculation, this serves to indicate the non-hydrodynamic scaling nature of the radiation transport. Naturally, the alpha-heating energy deposition in the hotspot also does not scale hydrodynamically; intuitively this is clear from the positive feedback interaction between alpha-heating and the hotspot temperature.

Using the same models (without alpha-heating), the yield was also predicted to scale as $\sim E_L^{3/2}$, or equivalently $\sim S^{4.5}$ [128]. This sets a lower bound on the expected yield scaling from this process when including alpha-heating.

3 Alpha-Heating and Chimera

In this chapter we describe the development of a charged particle transport package for the radiation-magnetohydrodynamics code Chimera, with the primary application of modelling fast fusion alpha-particles during the ignition and burn process in ICF experiments. Section 3.1 describes Chimera and the physics currently implemented. Section 3.2 then explains the potential ways to approach modelling fast alpha-particles, and Section 3.3 gives details on the implementation of the Coulomb collision framework and the relevant algorithms. Section 3.4 clarifies the underlying computational structures necessary to implement these algorithms. The unit and regression tests of the implementation are then described in section 3.5.

3.1 Chimera

All simulations described in this thesis were performed using the radiation-magneto-hydrodynamics code Chimera. Chimera has been developed over a number of years by various contributors, and undergone numerous upgrades during this time. Recent upgrades include the conversion of the radiation transport from single-group to multi-group [130], and the implementation of extended magnetohydrodynamics (MHD) terms [131].

Chimera uses a 3D Eulerian mesh in any of Cartesian (x, y, z), cylindrical (r, z, θ) or spherical (r, θ, ϕ) geometries. Domain decomposition is used to split the computational load, such that each processor simulates only a subset of the total simulation domain. The MPI protocol is used to communicate boundary information between processors.

3.1.1 Hydrodynamics

In the case of no MHD (as applies to this work), Chimera solves the following fluid equations:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{v}) = 0 \tag{3.1a}$$

$$\frac{\partial \rho \mathbf{v}}{\partial t} + \nabla \cdot (\rho \mathbf{v} \mathbf{v}) = -\nabla (P_i + P_e) + F_\alpha \tag{3.1b}$$

$$\frac{\partial \epsilon_e}{\partial t} + \nabla \cdot \epsilon_e \mathbf{v} = -P_e \nabla \cdot \mathbf{v} + \nabla \cdot (\kappa_e \nabla T_e) + Q_{ei} + Q_\gamma + Q_{\alpha,e} \tag{3.1c}$$

$$\frac{\partial \epsilon_i}{\partial t} + \nabla \cdot \epsilon_i \mathbf{v} = -P_i \nabla \cdot \mathbf{v} + \nabla \cdot (\kappa_i \nabla T_i) - Q_{ei} + Q_{\alpha,i} \tag{3.1d}$$

for a fluid of density ρ and velocity $\mathbf{v}$. Electrons and ions are treated as separate species, with individual energy densities $\epsilon_{e,i}$, temperatures $T_{e,i}$, conductivities $\kappa_{e,i}$ and pressures $P_{e,i}$. Quasineutrality is assumed, such that the electron and ion densities have the same velocity, and therefore they are considered as a single fluid for the hydrodynamics. F_α is the effect of the alpha-particle Coulomb collisions on the momenta of the electron and ion fluids. Q_{ei} is the collisional energy exchange between the electrons and ions, Q_γ is the energy exchanged between the electrons and the radiation field, and $Q_{\alpha,e/i}$ is the energy deposited into each species due to Coulomb collisions with alpha-particles. As explained in section 2.2.4, the neutron contribution to the heating of the hotspot is negligible and is therefore ignored.

The hydrodynamics solver consists of a 2nd order van Leer advection [132, 133], in combination with a von Neumann-Richtmeyer artificial viscosity. A simple line interface calculation (SLIC) [134] based method is used to maintain an approximate interface between the ablator and fuel materials to advect the materials separately.

3.1.2 Equation of State

Chimera uses tabulated equation of state data calculated using the Frankfurt equation of state (FEoS) model [135–137], with a separate table for each material. The thermodynamic variables are calculated from the Helmholtz free energy, F:

$$F = F_e + F_i + F_{bond}$$

where F_e and F_i are the electronic and ionic contributions respectively, and F_{bond} is a correction to the bonding to achieve zero pressure at solid density. The electron contribution is calculated using Thomas-Fermi theory, i.e. assuming a Fermi gas in the presence of an atom. The ionic contribution is calculated using the Cowan Model, which interpolates between various empirical models valid over smaller subsets of the density and temperature ranges. These tables are pre-calculated and tabulated offline, with an inline look-up procedure in order to expedite calculations. The look-up procedure uses the density and internal energy of each material to look-up electron and ion temperatures, which are then used in conjunction with the densities to find the pressure, sound speed and ionisation.

3.1.3 Thermal Conduction

The thermal conduction heat flow in Chimera, as seen above in the hydrodynamic equations, is given by:

$$\nabla \cdot q = \nabla \cdot (\kappa \nabla T) \tag{3.2}$$

where the thermal conductivity, κ, is calculated following Spitzer and Braginskii [138]:

$$\kappa_e = A_{T,e} \frac{n_e T_e^{5/2}}{(\sum_j Z_j^2 n_j) \ln \Lambda} \tag{3.3}$$

$$\kappa_i = A_{T,i} \frac{n_i T_i^{5/2}}{Z_{eff}^2 (\sum_j Z_j^2 n_j) \ln \Lambda} \tag{3.4}$$

where the coefficients $A_{T,e/i}$ also include the Epperlein and Haines [114] modifications in the case of a magnetised plasma, $\sum_j$ are sums over different material species, j, Z_j is the effective ionisation of species, j (i.e. allowing for partial ionisation) and Z_{eff} is the overall effective ionisation of the plasma.

The thermal conduction algorithms were recently upgraded by C. A. Walsh [131] (along with extended MHD capabilities, not applicable to this work) to include a super-stepping scheme [139]. This scheme reduces the number of subcycles, s, required in the thermal conduction algorithm (from s to $\sqrt{s}$), and in doing so speeds up calculations.

3.1.4 Radiation Transport

Chimera solves the following equations for radiation of energy flux, $\mathbf{F}_\nu$, and energy density, U_ν, at frequency, ν, propagating through material of emissivity, j_ν, and opacity, k_ν:

$$\frac{\partial U_\nu}{\partial t} + \nabla \cdot \mathbf{F}_\nu = 4\pi j_\nu - k_\nu c U_\nu \tag{3.5}$$

$$\frac{1}{c} \frac{\partial \mathbf{F}_\nu}{\partial t} + c\nabla \cdot \mathsf{P}_\nu = -k_\nu \mathbf{F}_\nu \tag{3.6}$$

Equation 3.5 represents the conservation of energy for radiation of a given frequency, and equation 3.6 the conservation of momentum. The equations are closed using the P_1 approximation [140] — which substitutes $\mathsf{P}_\nu = \frac{1}{3} U_\nu \mathsf{I}$ into equation 3.6 — giving non-diffusive transport. This is then adjusted into $P_{1/3}$ to obtain the correct wave-speed in the free-streaming limit (see A.3 for more details):

$$\frac{1}{3c} \frac{\partial \mathbf{F}_\nu}{\partial t} + \frac{c}{3} \nabla U_\nu = -k_\nu \mathbf{F}_\nu \tag{3.7}$$

The equations are solved for multiple radiation groups. Automatic flux limiting [141] analytically integrates equations 3.5 and 3.6 with an integrating factor $\exp(ck_\nu dt)$, and varies the flux smoothly between the diffusive and free-streaming limits. In doing so, this also prevents the transport of more energy than the free-streaming limit, which is possible under the $P_{1/3}$ approximation (but not physical). This method provides smooth flux variation between the diffusive and free-streaming limits. Here, the radiation transport interacts with the hydrodynamics through the change in energy density of the fluid, Q_γ (in equation 3.1). This is calculated using the divergence of the flux, $\int \nabla \cdot \mathbf{F}_\nu d\nu$, as well the direct emission and absorption contributions from the plasma. A more comprehensive description of the algorithms is given by C. Jennings [142] and K. McGlinchey [130].

3.1.4.1 Atomic Data

The detailed configuration accounting (DCA) code Spk is used to provide the atomic data, developed at Imperial College London by N. Niasse [143]. Spk uses the screened hydrogenic model (SHM) energy levels, including nl splitting [144], and supplemented with experimental energy levels from the NIST atomic database [145]. The population of these energy levels is solved for using a modified Saha equation following an effective temperature method [146], which allows for coronal-radiative (coronal) equilibrium (CRE); CRE allows for equilibrium between radiative de-excitation with collisional excitation and de-excitation. The opacities and emissivities are calculated and tabulated offline, and then looked up in-line. Fully resolved spectra at important conditions which the materials will undergo are used to inform the choice of frequency groups, striving to capture as much of the spectral detail as possible [130], such as ionisation edges and line structure.

3.1.5 Fusion Products

We refer the reader to section 3.2 below, which discusses the various approaches to modelling fusion alpha-particles, including those implemented into Chimera. Neutron deposition is assumed negligible.

3.2 Approaches to alpha-particle modelling

There are several different approaches to modelling fast-fusion alpha-particles, each with varying degrees of accuracy and ease of implementation.

The simplest is local instantaneous heating, which considers only the amount of energy contained in the alpha-particles, and ignores any transport of said energy. This is the method originally implemented in Chimera, as it is straightforward to implement, wherein the number of fusion reactions within a given cell is used to calculate the total amount of alpha-particle energy released. This energy is partitioned among the electron and ion fluids within the cell in accordance with the strength of the interactions between alpha-particles and the respective fluids (see sections 2.2.4 for further explanation). This partitioning can be approximated [12] to give equation 2.35:

$$f_i = \frac{1}{1 + \frac{32}{T_e}}$$

for the fraction of α-energy deposited into the ion fluid.

One level up from this is local, non-instantaneous (or delayed) heating. This accounts for the finite time over which the slowing via Coulomb collisions occur by calculating the energy as in the local instantaneous case, and depositing it in the local cell over the

finite timescale (adjusted from equation 2.36):

$$\tau_{slow} = \frac{1.3 \times 10^3 T_e^{\frac{3}{2}}}{\rho \ln \Lambda_{\alpha e}} ps \qquad (3.8)$$

where T_e is in eV and ρ is in kgm^{-3}.

However, the failure of both of these models to account for the finite alpha-range leads to significant inaccuracies. Local deposition results in the fusion heating only regions of plasma which are already the hottest, leading to significantly faster pressure build-up in a smaller region than would result from the non-localised heating of a much larger volume. The energy transport out of the hotspot is restricted primarily to thermal conduction (and radiation transport), with no direct alpha transport of energy, and therefore naturally produces much hotter hotspots. Clearly such models cannot be used to accurately study burn propagation, but are best-suited for order-of-magnitude calculations of experiments with low levels of alpha-heating.

It is much more physically accurate to model the non-locality of the alpha-deposition due to the high velocity of the alpha-particles. This charged particle transport through plasma can be done using deterministic or stochastic (kinetic) models. Deterministic models attempt to solve the Fokker-Planck equation, for example by reducing it to a diffusion equation [147], or by using discrete ordinate methods [148–150]. In these models, the energy- (or velocity-) dependence of the Coulomb scattering can be treated using a multi-group structure. Diffusive models are based on an assumption of isotropy, and thus struggle with anisotropic sources and sharp features, which become washed out by the diffusive transport. Discrete ordinates methods are best suited for highly anisotropic sources, or 1D simulations where the isotropy is given by the geometry. In multiple dimensions, the ray effect [151] arises naturally from the discretization of angles - that is, quantities are transported preferentially along particular angular directions, which is non-physical.

Kinetic (or stochastic/Monte-Carlo) modelling of alpha-particles is the most physically accurate with the fewest approximations involved, but is also the most computationally demanding, requiring large numbers of particles to maximise statistical and physical accuracy, in addition to large quantities of random numbers and particle memory. There are two general approaches for implementing Coulomb collisions in particle simulations in plasmas; the binary approach and the grid-based Langevin approach. In the binary approach, particles in a cell undergo Coulomb collisions in pairs, with particles paired off and scattered off one another, conserving energy and momentum in the process [152, 153]. This approach is suited to modelling multiple species as particles, such as in full Particle-in-Cell (PIC) simulations of plasmas with both electrons and ions modelled as particles. The grid-based approach reduces the collisional Fokker-Planck equation to Langevin form [154, 155], allowing the calculation of a force on the particles using grid-based properties. This is well-suited for hybrid PIC-fluid codes, since the force can be derived from background fluid variables.

More recently, a third hybrid Vlasov-Fokker-Planck (VFP) approach has been developed to kinetically alpha-particles, rather than using a PIC approach [156, 157]. In this hybrid approach, the VFP equation for the particle distribution function is solved in 1D for both the background thermal DT ions and the suprathermal alpha-particles, while the electrons are solved as a fluid. The treatment of the thermalisation process via the distribution functions is more precise than the use of a thermalisation threshold energy, as in many particle-based approaches. This formalism also includes the effect of electrostatic fields generated by the pressure gradient, but is only computationally tractable in 1D.

3.3 Implementation

In this section, we describe the Coulomb collision framework for the particle push, as well as the algorithms for the particle spawn and population management.

3.3.1 Particle Push

We follow the framework set out by Sherlock [158] for implementing Coulomb collisions between particles incident on a Maxwellian background fluid moving at a finite velocity. The particle experiences both a deterministic frictional force and a stochastic velocity-space diffusion due to scattering, with the coefficients determined by the fluid properties and the particle velocity. For a particle α of charge $Z_\alpha e$, mass m_α and velocity v_α scattering off background fluid β of charge $Z_\beta e$, mass m_β, number density n_β, temperature T_β and thermal velocity $v_\beta = \sqrt{2k_B T_\beta / m_\beta}$, the slowing and diffusive coefficients [158] in the Spitzer formulation are given (repeated from equations 2.4-2.6 in section 2.1.1):

$$\frac{\partial v_{\alpha\|}}{\partial t} = -A(1 + m_\alpha/m_\beta)\frac{G(w)}{v_\beta^2}$$

$$\frac{\partial v_{\alpha\|}^2}{\partial t} = AG(w)/v_\alpha$$

$$\frac{\partial v_{\alpha\perp}^2}{\partial t} = A\frac{\Phi(w) - G(w)}{v_\alpha}$$

$$A = \frac{Z_\alpha^2 Z_\beta^2 e^4 n_\beta \ln \Lambda_{\alpha\beta}}{2\pi m_\alpha^2 \epsilon_0^2}$$

$$G(w) = \frac{\mathrm{erf}(w) - w\frac{\partial \mathrm{erf}(w)}{\partial w}}{2w^2}$$

where $w = v_\alpha/v_\beta$, $\mathrm{erf}(w)$ is the error function and $G(w)$ is the Chandresekhar function.

The finite background fluid velocity is accounted for by transforming into the co-moving frame before calculating the coefficients for particle slowing. Equation 2.4 can be used straightforwardly to calculate the frictional contribution to the change in particle velocity

$$\Delta v_{\|,fric} = \Delta t \frac{\partial v_{\alpha\|}}{\partial t} \tag{3.9}$$

while equations 2.5 and 2.6 are used in calculating standard deviations $\sigma_\parallel = \sqrt{\Delta t \partial_t v_\parallel^2}$ and $\sigma_\perp = \sqrt{\Delta t \partial_t v_\perp^2}$. The parallel diffusion contribution is given simply by

$$\Delta v_{\parallel, diff} = \hat{N}(\sigma_\parallel) \tag{3.10}$$

where $\hat{N}(\sigma)$ is the selection of a random number from a Gaussian distribution of standard deviation σ. The perpendicular diffusion contributions in orthogonal directions (i, j) are given by

$$\Delta v_{\perp, \{i,j\}} = \hat{N}(\sigma_\perp) \times \{\cos(\theta_\perp), \sin(\theta_\perp)\} \tag{3.11}$$

for random angle $\theta_\perp \in [0, 2\pi)$.

Particle motion is integrated using a leapfrog scheme, in which the position and velocity updates are half-timesteps out of sync with one another:

$$\mathbf{x}_{i+1} = \mathbf{x}_i + \mathbf{v}_{i+\frac{1}{2}}\Delta t \tag{3.12}$$

$$\mathbf{a}_i = F(\mathbf{x}_i) \tag{3.13}$$

$$\mathbf{v}_{i+\frac{1}{2}} = \mathbf{v}_{i-\frac{1}{2}} + \mathbf{a}_i \Delta t \tag{3.14}$$

where the x_i is the position at timestep $t = t_i$, $v_{i+\frac{1}{2}}$ is the velocity at half-timestep $t = t_{i+\frac{1}{2}}$, and a_i is the acceleration at timestep $t = t_i$. This integration algorithm is accurate to second order in time (i.e. errors scale $\sim (\Delta t)^2$) [159], and is subcycled according to the particle cell-crossing time and the slowing (τ_s) and scattering ($\tau_\parallel, \tau_\perp$) relaxation times:

$$\tau_s = \frac{v_\alpha}{\partial v_\alpha / \partial t}, \; \tau_\parallel = \frac{v_\alpha^2}{\partial v_{\alpha\parallel}^2 / \partial t}, \; \tau_\perp = \frac{v_\alpha^2}{\partial v_{\alpha\perp}^2 / \partial t} \tag{3.15}$$

Note that the particles are slowed against both ions and electrons in our scheme, and therefore each coefficient is calculated twice.

The framework allows for exact conservation of energy and momentum. The conservation of momenta from individual particles crossing the grid cell can be used to calculate the change in fluid momentum and thus the change in fluid kinetic energy (ΔK_{fluid}), and the conservation of energy used to calculate the change in fluid energy (ΔE_{fluid}). This then gives the change in thermal energy in the fluid, $\Delta U_{fluid} = \Delta E_{fluid} - \Delta K_{fluid}$. Our calculations are expedited using approximate forms [160] for particle-electron and particle-ion Coulomb logarithms, in addition to the Sherlock's approximation for $G(x)$ [158]:

$$G(v_\alpha / v_\beta) \approx \frac{v_\alpha v_\beta^2}{2 v_\alpha^3 + \frac{3\sqrt{\pi}}{2} v_\beta^3} \tag{3.16}$$

Magnetised transport

The scheme can be extended to magnetised transport of the charged particles, where the Lorentz equation can be rearranged to give:

$$\frac{\mathbf{v}_{t+\Delta t/2} - \mathbf{v}_{t-\Delta t/2}}{\Delta t} = \frac{q}{m}\left[\mathbf{E} + \frac{\mathbf{v}_{t+\Delta t/2} + \mathbf{v}_{t-\Delta t/2}}{2} \times \mathbf{B}\right] \tag{3.17}$$

Following the Boris method [161], we define:

$$\mathbf{v}^{\pm} = \mathbf{v}_{t\pm\Delta t/2} \mp \frac{q\mathbf{E}}{m}\frac{\Delta t}{2} \tag{3.18}$$

such that equation 3.17 becomes:

$$\frac{\mathbf{v}_{t+\Delta t/2} - \mathbf{v}_{t-\Delta t/2}}{\Delta t} = \frac{q}{2m}\left(\mathbf{v}^{+} + \mathbf{v}^{-}\right) \times \mathbf{B} \tag{3.19}$$

resulting in a simple $\mathbf{v} \times \mathbf{B}$ rotation. Ignoring the contribution from the electric field E (which is already included through the stopping model), then $\mathbf{v}^{\pm} = \mathbf{v}_{t\pm\Delta t/2}$. The $\mathbf{v} \times \mathbf{B}$ rotation through an angle θ can be achieved through an intermediate vector, $\mathbf{v}'$, which lies at the halfway angle of rotation $\theta/2$ to $\mathbf{v}^{-}$, giving:

$$\mathbf{v}' = \mathbf{v}^{-} + \mathbf{v}^{-} \times \mathbf{t} \tag{3.20}$$

where

$$\mathbf{t} \equiv -\hat{\mathbf{b}}\tan\frac{\theta}{2} = \frac{Z_{\alpha}e\mathbf{B}}{m}\frac{\Delta t}{2} \tag{3.21}$$

is obtained through geometry. The final part of the rotation is given by:

$$\mathbf{v}^{+} = \mathbf{v}^{-} + \mathbf{v}' \times \mathbf{s} \tag{3.22}$$

where again, by geometry:

$$\mathbf{s} \equiv -\hat{\mathbf{b}}\sin\theta = \frac{2\mathbf{t}}{1+t^2} \tag{3.23}$$

This rotation scheme sets an additional constraint on the time-step, in order to accurately resolve the particle Larmor radius, $r_g = m_{\alpha}v_{\perp}/e|Z_{\alpha}||\mathbf{B}|$. With a gyroperiod of $T_g = 2\pi r_g/v_{\perp} = 2\pi m_{\alpha}/e|Z_{\alpha}||\mathbf{B}|$, an angular rotation limit of $\Delta\theta_{mag}$ gives a magnetic time-step of:

$$\Delta t_{mag} = \Delta\theta_{mag}m_{\alpha}/eZ_{\alpha}B \tag{3.24}$$

for a particle α of charge Z_{α} and mass m_{α}.

In testing this, we found that an angular limit of $\Delta\theta_{mag} = 1°$ was stable but not necessarily 100% accurate in maintaining adiabatic invariants in a dipole field, but the $\Delta\theta_{mag} = 0.1°$ required for accurate resolution was prohibitively small, resulting in significant increases in computational runtime.

This work does not include the effects of magnetic fields, and as such this magnetised transport is not used in this thesis. However, it is included here for completeness, and

we note that the magnetised transport has been used in other work in the group.

3.3.2 Particle Spawn

Since for our model we are only interested in the nuclear reaction involving $D + T \rightarrow \alpha\left(3.5 MeV\right) + n\left(14.1 MeV\right)$, this allows us to spawn alpha-particles in only the regions containing significant amounts of fuel. We use the Bosch-Hale formulae [162] to calculate the reactivity $\langle\sigma\nu\rangle_{DT}$, and then calculate the number of reactions in the cell as:

$$N_{reactions} = \frac{1}{4}n_{DT}^2 \langle\sigma\nu\rangle_{DT} \Delta V \Delta t \tag{3.25}$$

for an ion fluid of equimolar DT within grid cell of volume ΔV in a time-step of Δt. In order to keep the computational load small, we impose a constant minimum threshold on the number of reactions per cell required to spawn macro-particles, N_{min}, which can be set at the initialisation of the simulation. If $N_{reactions} > N_{min}$, then the computational weighting for macro-particles to be spawned is set as $N_\alpha = N_{reactions}/N_{cell}$ where N_{cell} is the number of computational particles to spawn per cell, and is a constant set at initialisation.

Macro-particles are spawned at random locations within each cell with an energy distribution broadened [163] around the $3.5 MeV$ by the reacting plasma distribution, and with a spherically isotropic velocity distribution. The cell fluid velocity is then added to the macro-particle velocity, to transform from the fluid rest frame to the lab frame.

The reacting DT particle pairs will also have much higher energies than the average DT pair, and thus fusion reactions will remove particles from the tail of the Maxwellian. Assuming that the repopulation of the tails is fast with respect to hydrodynamic and particle-transport timescales, we therefore use the mean energy of the reacting pair $E_0 = 66.46 T_i^{\frac{2}{3}}$ (for E_0, T_i in eV) [164] to calculate the loss in energy density from the reacting cell as:

$$\Delta U_{fusion} = -66.46 T_i^{\frac{2}{3}} \cdot e \cdot \frac{N_{reactions}}{\Delta V} \tag{3.26}$$

The reacting cell density is also reduced accordingly by:

$$\Delta\rho_{fusion} = \frac{1}{2}\frac{N_{reactions}}{\Delta V}\left(m_T + m_D\right) \tag{3.27}$$

3.3.3 Population Management

We employ an intuitive population management scheme in our model, based on the scheme employed in a PIC model previously developed at Imperial College London by S. Vickers [165] for MHD modelling in Z-pinches. We expect the macro-particles with the highest total energy $E_{total} = w_{macro} \times E_{macro}$ (where a macro-particle represents w_{macro} particles of energy E_{macro}) to have the most impact on the simulation. Hence we look to reduce the population by targetting those macro-particles with the lowest E_{total}.

We first dynamically calculate an energy threshold $E_{macro} = E_{thresh}$ below which some fraction of particles, $1 - f_{surv}$, will be removed. We do this by using the maximum value of $E_{total,max}$ in the macro-particle population of size N_{total}, such that $E_{thresh} = k E_{total,max}$ for some constant k. k is increased from a default value of 0.1 such that $f_{surv} \in (0,1)$, for:

$$f_{surv} = \frac{(N_{target} - N_{save})}{N_{PopControl}} \qquad (3.28)$$

where N_{target} is the target total population, N_{save} is the number of macro-particles with $E_{total} > E_{thresh}$, $N_{PopControl}$ the number of macro-particles with $E_{tot} \leq E_{thresh}$, and $N_{save} + N_{PopControl} = N_{total}$.

In other words, k is adjusted in order to ensure that f_{surv} falls within the range $(0,1)$ and sets the value for E_{thresh}, with typical k values of $\sim 0.1 - 0.2$. Once k and E_{thresh} are found, the population $N_{PopControl}$ is then reduced by a fraction $1 - f_{surv}$ through random annihilation of particles. To conserve particle number, the w_{macro} (and therefore E_{total}) of surviving f_{surv} macro-particles (within the population $N_{PopControl}$) are increased by a factor of $1/f_{surv}$.

In effect, macro-particles within the lowest total energy population are removed from the simulation, with the weights and energies of the surviving fraction of this population increased to conserve the total number of real particles.

3.4 Computational Methods

In this section, a description is given of the more fundamental computational methods, including random number generation and the data structures that are used to store the particle data.

3.4.1 Random Number Generation

For a Monte-Carlo simulation such as ours, the choice of random number generator (RNG) is of crucial importance, given the frequency with which it will be called. Ideally, the RNG would be both as fast and as random (uncorrelated) as possible. Although there are two branches of RNGs — true and pseudo — the focus of our discussion will remain with the pseudo-random number generators (PRNGs). A true random number generator (TRNG) typically relies on an unpredictable physical process, such as thermal noise or atmospheric radio noise in order to seed and generate random numbers. These numbers are truly random — but the generator is often very slow, and TRNGs are generally reserved for applications where the randomness is of critical importance, such as encryption. Here we give only a brief overview of some of the methods of generating pseudo-random numbers and the subsequent grounds for our choice of PRNG, while reserving a more in-depth summary of the field for elsewhere (for example, see [166, 167]).

PRNGs are RNGs which rely on algorithms to generate a deterministic sequence of

"pseudo-random" numbers rather than truly random numbers. However, these numbers can be generated quickly, do not require special hardware (e.g. to source the physical process), and due to the deterministic nature, allows one to generate the exact sequence of pseudo-random numbers repeatedly if needed (e.g. debugging). The algorithm used by the PRNG is key to the success or failure of the PRNG — a poor algorithm might produce streams of numbers with hard-to-detect underlying correlations. In general, a PRNG produces a sequence of numbers based on recurrence:

$$x_i = f(x_{i-1}, x_{i-2}, ..., x_{i-n}) \tag{3.29}$$

where n seed numbers are used to initiate the PRNG, and the function f is the heart of the PRNG which determines both the speed and the randomness of the PRNG [168]. We shall see that there are several properties of PRNGs which need to considered, namely: speed, statistical quality, period and parallelisability. It is worth mentioning some sage advice from Donald Knuth: "Random number generators should not be chosen at random" [166].

3.4.1.1 Speed

The speed of a PRNG is simple to assess — one need only look at the time taken to produce a random number stream of a given length. It is generally dependent on both the quantity and type of operations involved in the function f. Bitwise operations such as bitshifts or bitwise-XORs are significantly faster than mathematical operations such as modulo division.

Bitshifts involve shifting the actual bits of the memory in which the numbers are stored. For example, a left bitshift on an 8-bit integer 6 stored as 0000 0110 would give the integer 12, stored as 0000 1100.

XOR	$b = 0$	$b = 1$
$a = 0$	0	1
$a = 1$	1	0

Table 3.1: The truth table for the XOR operation for input bits a and b.

The bitwise-XOR involves applying the XOR operation to the bits in two numbers, with the truth table for XOR shown in table 3.1. In effect, the XOR operation between bits a and b produces 1 if and only if $a = 1$ OR $b = 1$, but produces 0 if $a = b$ (i.e. $a = 0, b = 0$ or $a = 1, b = 1$).

However, given the PRNG is unlikely to be the speed-limiting factor in the alpha-model calculations, it need only be fast "enough" to avoid this being the case. There is undoubtedly a trade-off to be made between speed and statistical quality - since one expects more operations to produce higher-quality streams (if said operations are selected carefully).

3.4.1.2 Statistical Quality

The "quality" of a PRNG is generally considered in terms of statistical correlations within the output stream of random numbers. Naturally, one expects PRNGs to produce streams of numbers which are as "random" as possible — that is, statistically uncorrelated. Therefore, correlations within a random number stream are an important indicator of the quality (or lack thereof) of a particular generator.

There are software suites designed for testing PRNG implementations, of which TestU01 [169] is the de-facto standard, containing a variety of statistical tests for both bit sequences and uniform random numbers over the interval $[0, 1)$. Passing these tests is necessary (but not sufficient) for a "good" PRNG.

3.4.1.3 Period

The period of a PRNG need only to be "long enough" — what this means in practise depends on the application at hand. Since the stream of (pseudo-)random numbers will repeat itself after the period of the PRNG, the period should be significantly longer than the expected number of calls to be made to the PRNG. For our purposes, a rough estimate requires a period of perhaps $\sim 10^{15}$ or higher, per processor. However, as long as the period satisfies these requirements, then larger periods no longer serve to indicate a better PRNG.

3.4.1.4 Parallelisation

In a code that is highly parallelised (such as Chimera), the PRNG must be able to produce p streams of numbers for p distinct processors, where the number streams are statistically uncorrelated both within the stream of each processor, and between streams of separate processors. There are several approaches to parallelisation of PRNGs, the most commonly used of which are "random seeding" and "parameterisation". Random seeding initiates the same PRNG in each processor with a different, "random" seed, in the hope that the resultant streams will be statistically independent of one another. However, there is no mathematical or theoretical basis that this will be true, and thus this is essentially choosing a random number generator at random.

Parameterisation changes the parameters in the PRNG for each processor, which is effectively changing f slightly. For example, if we use a linear congruential generator (LCG) (explained below in subsubsection 3.4.1.5):

$$x_{i+1} = (ax_i + c) \bmod m \tag{3.30}$$

to generate random number stream x_i, then parameterisation could vary the multiplier a or the constant c for each processor. However, this method has been demonstrated to have substantial statistical weaknesses [170], and finding sets of parameters (such as a, c, m for an LCG) which result in good, non-correlated streams for a particular algorithm is highly non-trivial. Therefore, this method should be avoided.

Two robust methods of parallelising a PRNG are "block splitting" and "leapfrogging" [171]. Block splitting involves splitting the stream x_i into blocks of length L such that each processor p_j receives the random number sequences $r_{j,i}$:

$$r_{0,i} = x_i$$
$$r_{1,i} = x_{i+L}$$
$$\vdots$$
$$r_{j,i} = x_{i+jL}$$

One must know (or be able to calculate) the length L of the block of numbers in advance, such that each processor only requires fewer than L random numbers. This method is feasible only if the underlying algorithm allows one to skip from x_i to x_{i+L} efficiently, without calculating all of the numbers in between.

Leapfrogging instead splits the random number stream x_i in the following manner over k processors for the n^{th} call to the PRNG:

$$r_{0,i} = x_{ni}$$
$$r_{1,i} = x_{ni+1}$$
$$\vdots$$
$$r_{k-1,i} = x_{ni+(k-1)}$$

This is the most versatile method, since no estimate of the number of numbers required by each processor is needed, only an algorithm which allows the efficient generation of only every k^{th} element in the sequence.

Unfortunately, requiring efficient block splitting or leapfrogging limits the possible PRNG choices significantly.

3.4.1.5 Potential choices and their bases

Linear congruential generators (LCGs) are the simplest and amongst the most common PRNGs, with a recurrence (repeated from above, equation 3.30):

$$x_{i+1} = (ax_i + c) \bmod m$$

to generate the random number stream x_i based on multiplier a, constant c and modulo divisor m. The quality and period of these generators are highly dependent on a, c, m and seed value x_0. While simple and fast, the output of LCGs will form a set of hyperplanes when plotted in n-dimensions. All possible values of the LCG will lie on these planes, with the number of planes dependent on the choice of a and m. In other words, plotting pairs of (x_n, x_{n-1}) will produce lines between the random number pairs in the 2D space, and similarly plotting (x_n, x_{n-1}, x_{n-2}) produces 2D planes in the 3D space. Figure 3.1 illustrates the spectral test for a simple LCG with $a = 7, m = 71$,
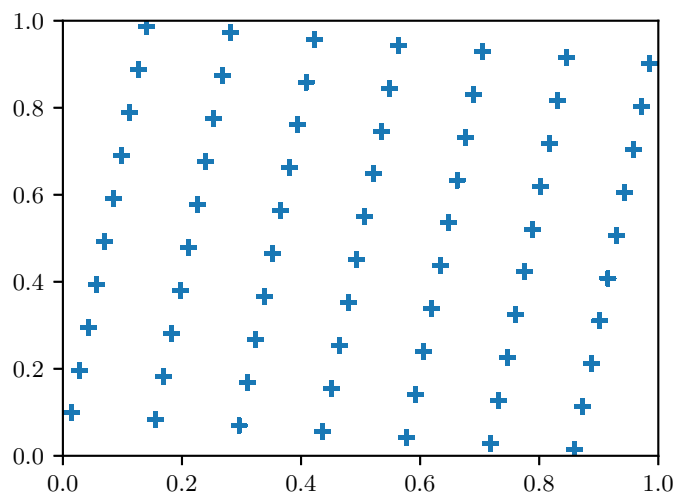
Figure 3.1: A plot of the spectral test showing clear correlations for a simple linear congruential generator with $a = 7$, $m = 71$.

showing the lines formed by plotting successive random number pairs. LCGs which can pass most tests can still fail this particular spectral test, which measures the distance between these planes.

Lagged Fibonacci generators are another class of PRNGs based on a generalisation of the Fibonacci sequence to give:

$$x_i = (x_{i-j} \odot x_{i-k}) \bmod m, \quad 0 < j < k \tag{3.33}$$

for binary operator $\odot$, such as addition or multiplication. These generators require a seeding block, and are highly dependent on the seed choice — a seed-block $x_k = 0 \, \forall \, k$ would result in a terrible PRNG!

The Mersenne Twister [172] is generalised feedback shift register (GFSR) PRNG, a class of Lagged Fibonacci generators which use the XOR operation. The algorithm has a period given by the Mersenne Prime $M_n = 2^n - 1, n \in \mathbb{N}$. The implementation `MT19937()` having $n = 19937$ is one of the most commonly used PRNGs, with implementations available as the default PRNG in many languages and libraries. However, in spite of its popularity, the Mersenne-Twister is no longer particularly fast or efficient, and is rather unwieldly algorithmically, requiring a large number of operations. The standard implementation `MT19937` requires hundreds of thousands of iterations to produce useful PRNGs, and it also requires a large amount of memory to store the internal and generator states, compared to other PRNGs. In addition, implementations with existing parallelisation capabilities are uncommon and slow.

The PRNG `ran2()` is another popular option from Numerical Recipes, and is commonly available in many languages by default. However, while it passes all statistical tests, it is slow and has a comparatively small period of $\sim 10^{18}$. It is also outdated and lacks parallelisability.

In choosing our PRNG, we consider not only the standard speed, quality, period and parallelisability requirements, but also ideally whether there is an existing implementation of the algorithm in Fortran — since most PRNGs are written in C++. This lead to the choice of PRNG based on the `xorshift` generators [173], which combine the bitwise-XOR operation with bitshifts (operations which are explained earlier in section 3.4.1.1). These are extremely fast and very simple — but also do not pass all statistical tests [174]. `xorshift` variants can be produced by combining the bitwise-XOR and bit-shift operations (explained above in section 3.4.1.1) with either addition (`xorshift+`) or multiplication (`xorshift*`) in 64-bit operations [175], which can eliminate the statistical problems while still maintaining the speed and simplicity. In fact, the `xorshift128+` generator proposed by Vigna [176] is the PRNG of choice in Google's Javascript engine V8, in addition to Firefox and Safari.

The generator `xorshift1024*` [175] was chosen as the default PRNG. With a period of $\sim 2^{1024}$, it is a high-quality generator with a fast, simple implementation (available also in Fortran as a translation from C++) and a low memory-footprint. There is also a `jump` function which performs the block splitting functionality described above, allowing for robust parallelisation. `jump` skips 2^{512} values, and therefore allows for 2^{512} strictly non-overlapping streams of length 2^{512}, which should cover both enough processors and individual calls within a processor. The (slightly) faster and more compact generator `xoroshiro128+` (the successor to xorshift128+ which also uses rotation operations) was also included, but is reserved for use with smaller-scale simulations (although with `jump` periods of $2^{64} \approx 10^{19}$, should still be sufficient for our purposes).

3.4.1.6 Non-uniform distributions - Gaussian

Most PRNGs produce uniformly distributed streams of bits, translated into either reals in interval $[0, 1)$ or integers between $[0, I_{max}]$ for maximum integer I_{max}. However, Gaussian-distributed random numbers are also needed, and there are a number of ways in which one can transform a uniform PRNG, U, producing a uniformly distributed number α in the range $(0, 1)$ to a Gaussian number, x. These can be reduced to 4 categories: cumulative density function (CDF) inversion, transformation, rejection and recursion.

The first of these simply inverts the CDF $\Phi(x)$ such that $x = \Phi^{-1}(\alpha)$. However, these require approximations for the Gaussian distribution, the quality of which affects the quality of the Gaussian random number generation. Transformation methods involve a direct transform from U to x. Rejection is like transformation, but also includes a conditional rejection dependent on the value of the transformation. The recursive method uses previously generated Gaussian numbers to produce new ones. Accurately sampling values from the tails of the distribution is particularly difficult, and for some methods requires a separate routine.

Thomas *et al.* [177] give an in-depth review of the available methods and algorithms for Gaussian random number generation, in addition to testing the quality and speed

of a significant number of these algorithms. They found that the Ziggurat method [178] was the second fastest, but did not suffer from correlations that the fastest algorithm suffered from — and therefore is the most appropriate.

We use the Ziggurat method as our default for generating Gaussian random numbers, using whichever uniform RNG is chosen for the simulation (by default `xorshift1024*`, but `xorshiro128+` also available, as detailed above).

3.4.2 Memory Structures

Since we will require large numbers of computational particles in order to maintain statistical accuracy, the question of how best to store the particle data is critical to the computational speed of the alpha-particle transport. We will first consider the demands that will be made of the particle data storage structure in our model, and then examine some common data structures to see how well they fit our requirements.

Requirements

At every hydrodynamic time-step, the number of new alpha-particles spawning throughout the hotspot will be on the order of 10^{14}. In an ideal world, we would have as many macro-particles as actual particles to obtain physical accuracy — however, this is simply infeasible with today's computational resources. The number of macro-particles and the physical weighting (the number of particles) represented by each macro-particle is therefore a compromise between computational expense and physical/statistical accuracy, and is affected by both the maximum cap on the number of macro-particles and the rate at which macro-particles are spawned.

Considering the hydrodynamic grid of Chimera, we are given the choice of either fixing the weight per macro-particle and dynamically varying the number of macro-particles spawned in each cell, or fixing the number of macro-particles per cell and varying the weighting of each macro-particle. Since we will be simulating capsules with yields from $\sim 10^{14} - 10^{18}$, the former option would result in variations in the number of macro-particles of ~ 5 orders of magnitude. The latter option provides much more predictable performance, with the number of macro-particles more easily estimated simulation to simulation. Regardless, the higher the total number of particles that can be simultaneously handled, the more accurate the simulation will be.

As an illustrative example, if we consider a hotspot of radius $30\mu m$ in a simulation of resolution $3\mu m$, then we are spawning particles in roughly $20^3 = 8000$ cells. Statistical accuracy would require that we need to be spawning somewhere in the region of $\sim 10^4 - 10^5$ macro-particles in any one spawn cycle, at a minimum.

In addition, alpha-particles with thermal velocities contribute very little to heating, and thus thermalised particle motion does not need to be calculated. Thermalised macro-particles should therefore be removed throughout the simulation, with a macro-particle removal rate similar to the rate of generation.

Thus, the primary demand on the data structure used to store our particle data is the ability to easily add and remove many elements (i.e. macro-particles, in this case) in any order, multiple times, and a secondary demand is being able to store a large number of particles. In addition, since the simulation domain is split across processors, we also need to be able to pass the information on the macro-particles between processors as they travel through the hotspot.

At each time step, all particles will need to be evolved individually, which means that we will need to cycle through all the particles at each time. Actions that need to be taken on an individual particle, such as thermalisation or transferring between simulation domains, will either need to be done during this cycle, or the particle marked in some manner, then found and processed in a call later on in the alpha-heating cycle.

3.4.2.1 Arrays

The simplest and most commonly used form of data storage is the array, consisting of a collection of indexed elements, typically stored in a contiguous section of memory. The index of an element in the array is used to identify the corresponding memory address for storage or selection, which can be achieved in constant time. Since the data is stored in contiguous memory, sequential access to elements is fast, and there is little memory-overhead, resulting in a relatively compact data structure.

Dynamic arrays allow for variable sizing of the array and therefore insertion and deletion of elements — by allocating an entire block of memory, and then using elements in this block (or not) as needed. However, if the allocated memory is not enough to store the entire dynamic array, then the entire array must be reallocated to a new block of memory, taking $\Theta(n)$ time for an array of n elements. This reallocation step also requires two chunks of contiguous memory, one for the existing (too small) array and one for the new (expanded) array. This can be seen in figure 3.2, showing the entire process of reallocation of dynamic arrays. Here, the new block of memory for the resized array large enough to contain the new elements to be inserted is allocated (1) and the existing elements then shifted across to the new block of memory (2). During these steps, the total memory required is the sum of that required for the existing array **and** the larger, resized array. After moving the existing elements, the old array can be deallocated and the memory requirement reduced to just the resized array (3), into which the new elements are then inserted (4).

Insertion/deletion of an element at any position also requires moving all the following elements of the array, which can also require expensive re-allocation if the array is not large enough. Moving all the elements can be avoided by storing a list of indices at which elements are deleted and into which new elements can be inserted, but such a list can rapidly become cumbersome with a large number of insertions and deletions. If one needs to search for a specific element in the array, this is fast if the index of the element is already known, but otherwise requires looping through the entire array.
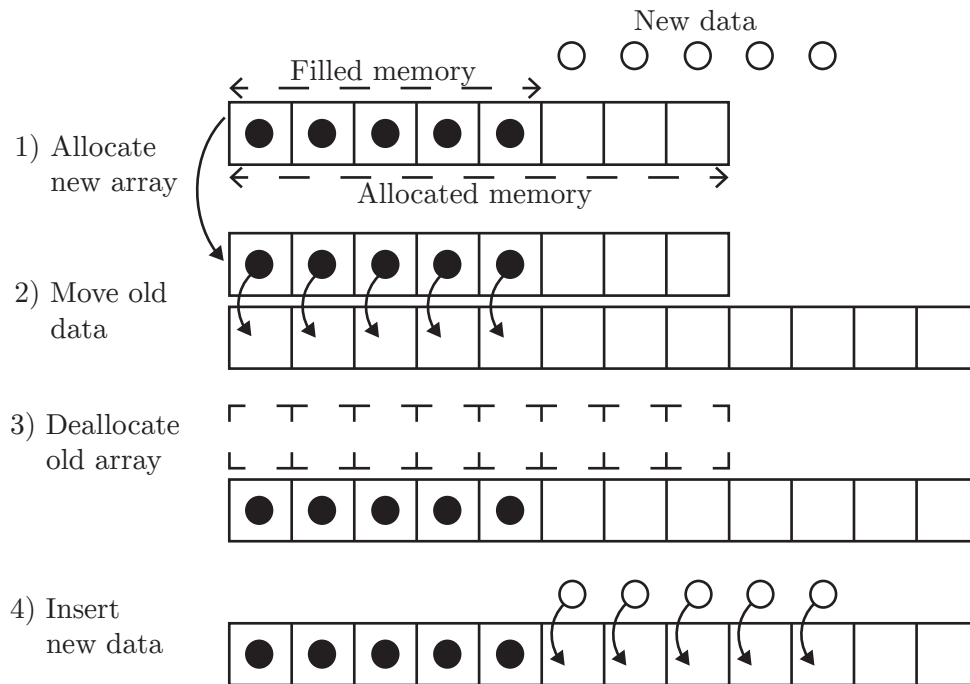
Figure 3.2: A schematic showing the process of reallocation of dynamic arrays to illustrate the memory overhead. 1) The memory required by the new data to be added to the initial array is larger than the available allocated memory, and so a new, larger array must be allocated. 2) The old data must then be moved between the two arrays, during which the total memory allocated is the memory required for both arrays. 3) Only once the old data is transferred can the old array be deallocated and the memory recovered. 4) The new data can be inserted into the new array, which has enough memory available to store it.

3.4.2.2 Linked Lists

Linked lists are a collection of nodes containing "pointers" and "data" fields, with the pointer field containing the memory address of the next node. Linking nodes in a linked list can be done in a number of ways, such as singly linked (in which each node links only to the next node), doubly linked (where each node links both backwards and forwards), circular (where the last node of the list links back to the first) and multiple (where each node links to two or more other nodes).

Naturally, linked lists have a significantly higher memory overhead per element than dynamic arrays do, since each element has references as well as data (as opposed to only one set of references for the entire structure). Elements in a linked list can only be accessed sequentially, and such sequential access is slower than that of dynamic arrays due to the incontiguous memory storage.

However, linked lists do not require a single contiguous block of memory for the entire structure, and therefore its length is limited not by the length of contiguous memory, but rather the total memory available. Use of a linked list also does not require *a priori* knowledge or estimates of the number of elements needed — the memory for each element can be allocated and deallocated at the point of insertion and deletion. Compared to the dynamic array, insertion/deletion operations are much faster due to avoiding the slow reallocation step, and removing the need to move all the following elements — one needs only to change the pointers of the surrounding nodes.

3.4.2.3 Comparison

Using dynamic arrays would not be computationally efficient, since we do not know *a priori* how many particles we will be spawning — indeed this will be changing every time-step with new particles spawning and thermalising. If a larger array is needed than initially allocated, we would need to allocate a new, larger block of continuous memory and move all the old data before deallocating the old memory — a highly computationally expensive task which becomes ever more problematic the larger the arrays involved (see Fig. 3.2). This structure puts an upper limit on the number of particles that the memory can handle as a function of available blocks of contiguous memory, which is more restrictive than the limit of the total memory available.

In addition, operations involving the deletion and insertion of new particles become even more complicated since the indices at which particles will be deleted or inserted will be scattered throughout the array — some variation of a stack would need to be implemented in order to track the indices of available memory within the array which can be reused before needing to reallocate. These insertions and deletions would be significantly slower than with linked lists, which allow rapid insertion and deletion of elements from the lists and therefore for easy spawning and thermalising of particles, without the restriction of contiguous memory blocks (provided the elements do not need sorting, which they do not in our scenario). This simple insertion and deletion of particles from lists, combined with the ease of merging lists, also enables easy handling of
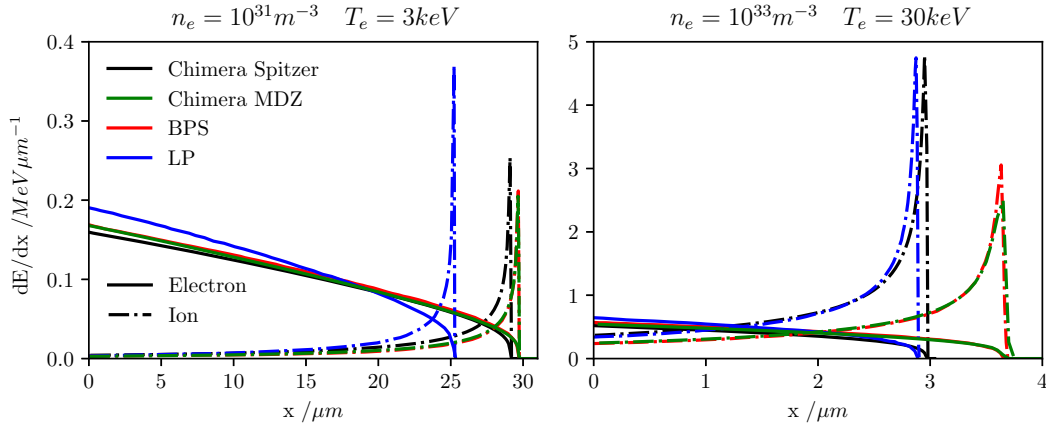
Figure 3.3: A comparison of Spitzer (black) and Maynard-Deutsch-Zimmerman (MDZ) (green) stopping powers from the Chimera burn module to the Li-Petrasso (LP) (blue) and Brown-Preston-Singleton (BPS) (red) models in a uniform background equimolar DT plasma at (left) $n_e = 10^{31}m^{-3}$ (equivalent to $\rho_{DT} = 4.1{\times}10^4 kgm^{-3}$) and $T_e = 3keV$, and (right) $n_e = 10^{33}m^{-3}$ and $T_e = 30keV$, following Singleton [103]. The stopping power is separated into contributions from electrons (solid) and ions (dash-dot).

any particles moving between processor domains. Therefore, we chose to use linked lists because of the more relaxed memory restrictions, as well as the speed and suitability of the method for implementing the processes required — namely spawn (insertion), thermalisation (deletion) and transfer between domains (moving particles between lists).

3.5 Testing

In this section, we describe the testing of the alpha model in order to ensure its accuracy. We compare the stopping power of our model to results from other models. The impact of various aspects of the model are tested in integrated simulations, including the spawn threshold, N_{min}, the population control and the overall convergence of yield with resolution. Finally, the ignition of an isobaric hotspot is examined as a comprehensive test problem.

3.5.1 Stopping Power

It is important that we are confident in the accuracy of the underlying microphysics of our charged-particle transport model, and therefore the stopping power of the model.

In figure 3.3, we compare the stopping power for the two models implemented in Chimera — Spitzer and Maynard-Deutsch-Zimmerman (MDZ) — with two established models: the Brown-Preston-Singleton (BPS) model [102] and the Li-Petrasso (LP) model [96]. The Chimera model stopping powers are calculated for $3.5MeV$ alpha-particles in a uniform background plasma of equimolar DT at: (**a**) electron number density, $n_e = 10^{31}m^{-3}$ and temperature $T_e = T_i = T = 3keV$, and (**b**) $n_e = 10^{33}m^{-3}$ and $T = 30keV$, comparing to values from Singleton [103] for BPS and LP. We find
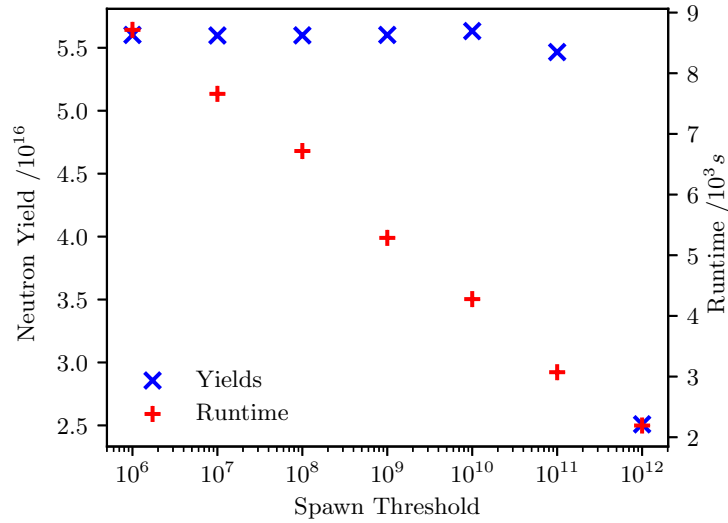
Figure 3.4: Yield (blue ×) and runtime (red +) as a function of the spawn threshold, N_{min}, (as defined in section 3.3.2) for integrated 1D simulations of N130927 at $1\mu m$ resolution.

good agreement of our models in stopping power and range for both electrons and ions with BPS, especially for the MDZ model, and in particular for the (**a**) conditions. The default stopping model used in the module is the MDZ model, owing to the excellent agreement with the BPS model, and the ease of calculation.

We note that throughout the majority of the slowing, the diffusive coefficients calculated using equations 2.5 and 2.6 are significantly smaller than the frictional, stopping power term of 2.4. For this reason, and the lack of other formulations with similar levels of computational simplicity, we calculate these diffusive coefficients using the Spitzer model, as in Sherlock and Rose [158] and section 3.3.1.

3.5.2 Spawn Threshold

In section 3.3.2, we defined a minimum threshold of reactions in a cell, N_{min}, which the number of reactions in the cell, $N_{reactions}$, was required to surpass (i.e. $N_{reactions} > N_{min}$) in order to spawn computational macro-particles within that cell. This reduces the number of particles being handled, and therefore the computational demand.

In figure 3.4, we examine the sensitivity of our model to this value within integrated 1D radial simulations of Highfoot shot N130927, restarted from peak radiation temperature (15.7ns) at $1\mu m$ resolution. (Details on how the simulations are set-up can be found in chapter 4.) The spawn threshold, N_{min}, was varied, and the neutron yield (blue ×) and simulation runtime (red +) examined. The yield is stable across the range of $N_{min} = 10^6 - 10^{11}$, dropping off only at $N_{min} = 10^{12}$.

We note that this value is particular to the simulation set-up. In general, N_{min} must be sufficiently low in order to be able to resolve the time-period of interest — in this case, the burn history. If we wish to resolve the burn pulse to within e.g. 1%, then
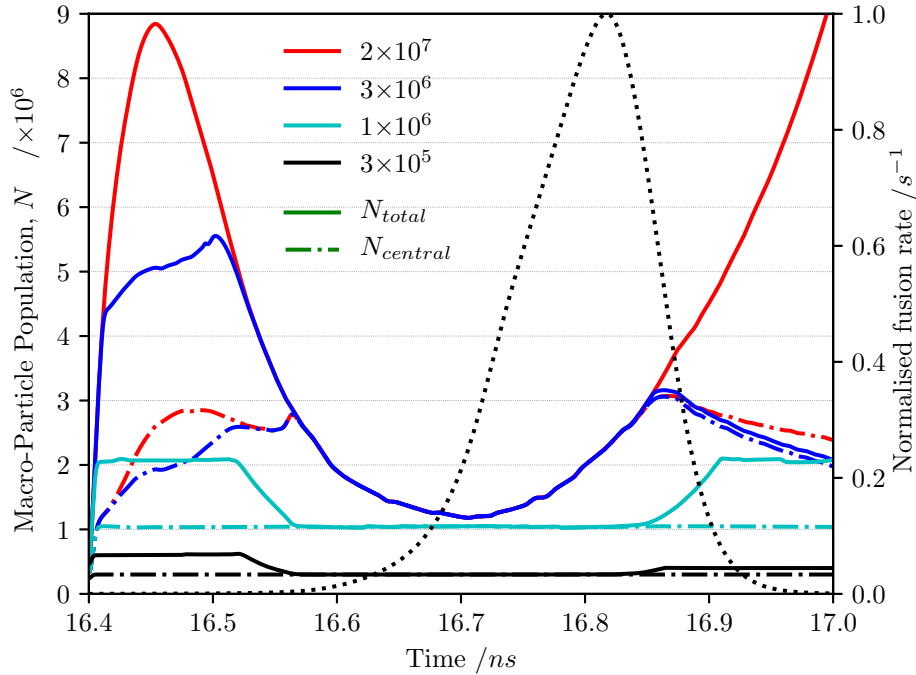
Figure 3.5: Variation in macro-particle number, N, over time, shown for all domains (N_{total}, solid) and only the central domain of the capsule ($N_{central}$, dash-dot) in 1D $0.5\mu m$ simulations of N130927 run with the following values of N_{target}: 2×10^7 (red), 3×10^6 (blue), 1×10^6 (cyan) and 3×10^5 (black). The dotted line indicates the fusion production rate (i.e. the burn history) for $N_{target} = 2 \times 10^7$, which is effectively unaffected by population control.

N_{min} must be below the production of individual cells at this point in time. Since the simulation has on the order of $\sim 10^2$ cells producing $N_{reactions} \sim 10^{12-14}$ fusion reactions per hydrodynamic time-step during the burn pulse (of width $\sim 100ps$), this gives a rough guideline of $N_{min} <\sim 10^{-2} \times O(N_{reactions})$.

3.5.3 Population Control

Section 3.3.3 described our population management scheme which targets and reduces the particle population by a proportion of the total target population, N_{target}. Due to the domain decomposition parallelisation, some domains (such as those containing a significant proportios of the hotspot) will produce significantly more macro-particles than other domains (e.g. regions containing very little fusing material). This is particularly prominent in 1D, where the simulations are split into radial domains, in which the central-most processors contain the majority of the macro-particles.

The effect of this can be seen in figure 3.5, which shows the variation in macro-particle number, N, over time for 1D simulations of N130927 run with $N_{target} = 2 \times 10^7$, 3×10^6, 1×10^6 and 3×10^5. These simulations are run at $0.5\mu m$ resolution and restarted at peak radiation temperature. Both the spawn threshold, N_{min}, and number of particles per cell, N_{cell}, (see section 3.3.2) were held constant at 10^7 and 1000 respectively. The
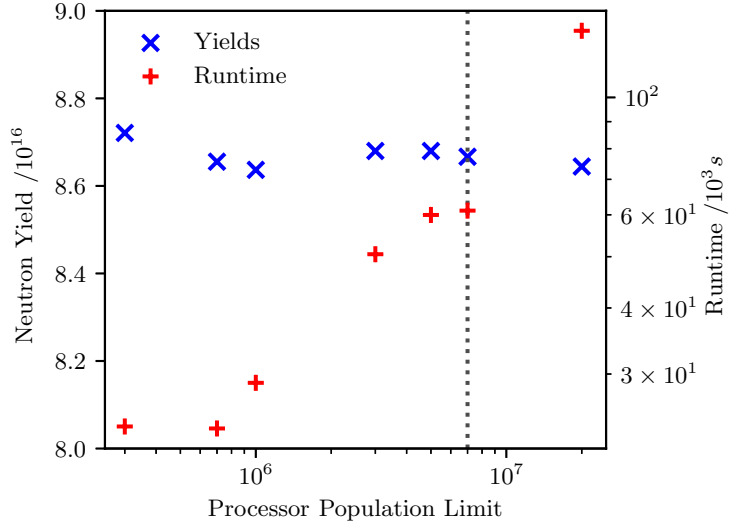
Figure 3.6: Yield (blue ×) and runtime (red +) as a function of the target total population per processor, N_{target} (as defined in section 3.3.3) for 1D simulations of N130927 at $0.5\mu m$ resolution. The dashed line indicates the maximum total number of macro-particles at any point for the uncapped case.

total number of macro-particles across all processors, N_{total}, and the number of macro-particles in just the most central domain, $N_{central}$, are both plotted. The central-most processor contains the centre of the hotspot, and therefore will contain the majority of the fusing particles during burn. The dotted line displays the normalised fusion rate for $N_{target} = 2 \times 10^7$, for which the simulation was effectively unaffected by population control as the macro-particle population did not hit the limit.

Considering first the uncapped simulation (red), it can be seen that after initialisation the total population, N_{total}, climbs to $\sim 9 \times 10^6$, of which $\sim 3 \times 10^6$ is in the central processor. As the capsule implodes, the region containing fusion compresses and N_{total} decreases as fusing material moves into the central processor, such that by $\sim 16.57ns$, only the central processor contains any fusing material and therefore any macro-particles. The macro-particle population variation with time is reminiscent of the time-variation of the hotspot radius, which decreases to a minimum, and then expands after peak compression as the hotspot begins to burn. Indeed, that is what is seen here, with N_{total} beginning to increase as the hotspot expands and a larger volume (and therefore more cells) producing fusion reactions.

Reducing the population limit to $N_{target} = 3 \times 10^6$ (blue) restricts N_{total}, such that the population immediately after initialisation is capped. N_{total} and $N_{central}$ vary in tandem up to $16.5ns$, indicating that the population is limited only in the outer processors rather than the central processor. As N_{total} falls, the $N_{target} = 3 \times 10^6$ simulation behaves very similarly to the uncapped simulation. A further reduction in the population limit to $N_{target} = 1 \times 10^6$ (cyan) results in all domains hitting the limit, indicated by flat lines for N_{total} and $N_{central}$. Even when there are macro-particles only in the central domain, N_{total} for $N_{target} = 1 \times 10^6$ is still below N_{total} for the uncapped population. This

is particularly important as it indicates that significant levels of population control are happening during the key ignition and burn phases of the implosion in which the majority of the fusion reactions are occurring. Additional reductions in the population limit to $N_{target} = 3 \times 10^5$ (black) result in similar capped behaviour.

Figure 3.6 examines the impact of this population limit, N_{target} on the yield (blue $\times$) and the runtime (red $+$). The dashed line indicates the maximum total number of particles in a particular processor at any given point ($\sim 7 \times 10^6$).

Comparing the results from figure 3.5 with figure 3.6, we see that the significant degree of population control in the smallest population limits actually has very little effect on the yield. Even as the population limit is set to $\sim 5\%$ of the uncapped maximum total number of particles, the simulation yield remains relatively constant, with less than 0.5% variation from the uncapped value.

Notice also that the degree to which population control is utilised does not appear to have a significant effect on the simulation runtime. The simulations with the most frequent population control are those with the lowest population limits, but these have the shortest runtime. The dominant effect appears to be that of a greater total number of particles increasing the simulation runtime.

3.5.4 Convergence

In this section, we examine the convergence of the simulation DT neutron yield with spatial resolution. Since the alpha-heating process results in a positive feedback loop, any small variations in burn-off yield might be enhanced through the alpha-heating process. In addition, high-resolution (small cell-size) simulations take significantly longer to run, particularly in 3D — a doubling in resolution increases the number of cells by 2^3, and the simulation time-step halves, resulting in a $16\times$ increase in runtime (not including the delay due to MPI communication between processors). Because of this, it is preferable to use as low a resolution as convergence will allow, and if not, to know how the simulations would be affected by smaller cell-sizes and therefore extrapolate.

In figure 3.7, DT neutron yield variation of 1D simulations of N130927 with increasing resolution is shown for both alpha-heating on (in blue $\times$) and off (in red $+$). All simulations were restarted at the same implosion time by merging cells from the same $0.25\mu m$ dataset, using a frequency-dependent radiation drive and explicit thermal conduction with sufficiently high subcycles to avoid hitting flux limiters.

We see that the alpha-heating off case converges steadily with decreasing cell-size, and by $0.5\mu m$ has converged to the $0.25\mu m$ yield. However, the alpha-heating on case displays a small step-change between $1.5\mu m$ and $1.0\mu m$, beyond which the yield continues to increase but at a much slower rate. The graph indicates that a cell-size of $1.0\mu m$ is mostly converged, with relatively small increases in yield beyond this resolution. Note that the positive feedback loop of alpha-heating necessarily stresses the other models in Chimera as well (such as thermal conduction, radiation transport, EoS). Therefore, this convergence test is not only a test of the alpha-model itself, but an integrated test
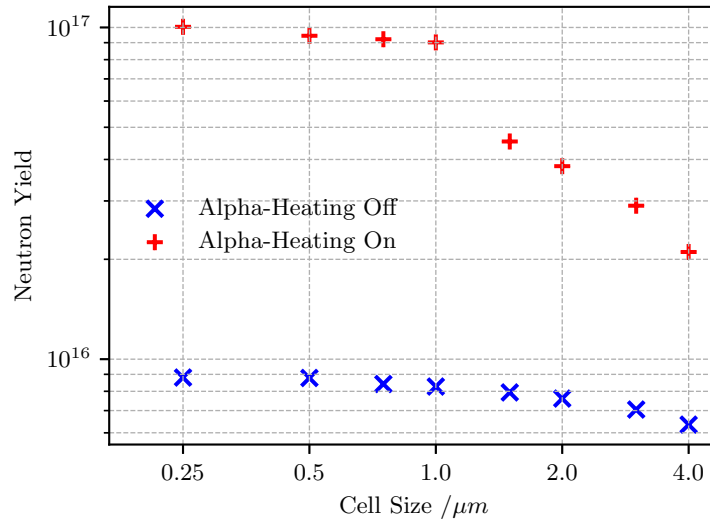
Figure 3.7: Yield against cell-size for alpha-heating on (blue ×) and off (red +). Yields appear to converge at cell sizes below $1\mu m$.

of the entire code applied to modelling ICF implosions with significant alpha-heating.

3.5.5 Isobaric Hotspot Evolution

Finally, we examine an isobaric hotspot ignition and burn problem with initial hotspot conditions of $\rho = 4{\times}10^4 kgm^{-3}$, $T_i = 7keV$ as a comprehensive test problem which further stress-tests the integration of the alpha-heating with other models in the code. Figure 3.8 shows the results of a 1D simulation by Atzeni and Meyer-ter-Vehn [7], with radial profiles of (a) ion temperature, T_i, (b) density, ρ and (c) pressure, at times (1) $t = 0ps$, (2) $t = 100ps$, (3) $t = 120ps$, (4) $t = 130ps$ and (5) $t = 140ps$.

We can compare this with the results from Chimera using both MDZ (left) and Spitzer (right) stopping models in figure 3.9, showing the temperature, density and pressure profiles as in figure 3.8, in addition to the alpha-heating power deposition (Q_α) profiles in the bottom panels (g,h). The initial profiles are slightly different, with a smoothed transition between hotspot and shell in order to avoid the generation of non-physical shocks due to discontinuities.

Comparing first the MDZ results to those of Atzeni, we note a similar time-evolution in temperature profiles, but with a sharper edge in the MDZ case. Similarly, the peaks in the density profile are also sharper. The peaks in the pressure profile coincide with the radial position of the peaks in the alpha-deposition, as well as with the drop-off in temperature and rise in density. Curves (2)-(4) in MDZ show the deflagration wave propagating through the fuel, with a deflagration-detonation transition [113] occurring between (4) and (5).

The (Chimera) Spitzer profiles are similar in their shape of evolution to those with the MDZ model, except that the evolution occurs faster for Spitzer — the deflagration-detonation transition appears to occur between (3) and (4) for Spitzer. This is due
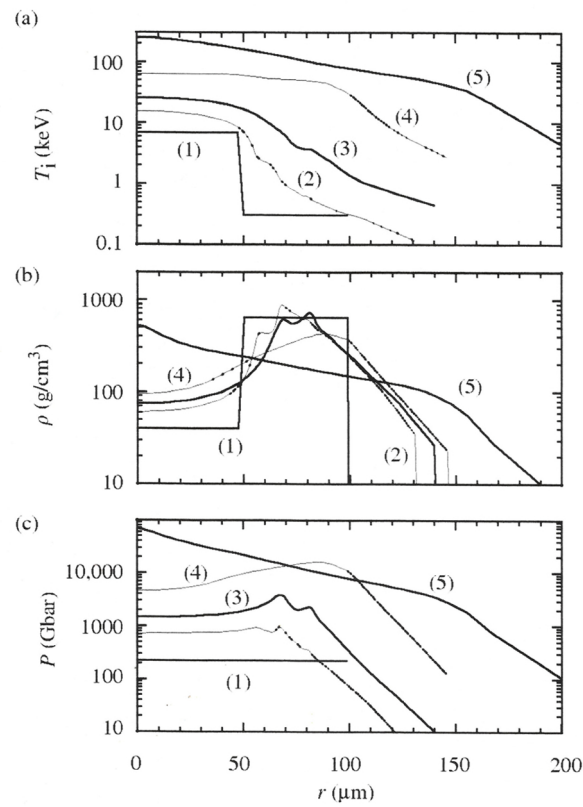
Figure 3.8: 1D radial profiles of a simulation of ignition and burn for an initially iso-baric equimolar DT capsule, showing (a) ion temperature, (b) density and (c) pressure at (1) $t = 0ps$, (2) $t = 100ps$, (3) $t = 120ps$, (4) $t = 130ps$ and (5) $t = 140ps$. Initial conditions for the hotspot are $\rho = 4 \times 10^4 kgm^{-3}$, $T_i = 7keV$. Reproduced from [7] with the permission of Oxford Publishing Limited.
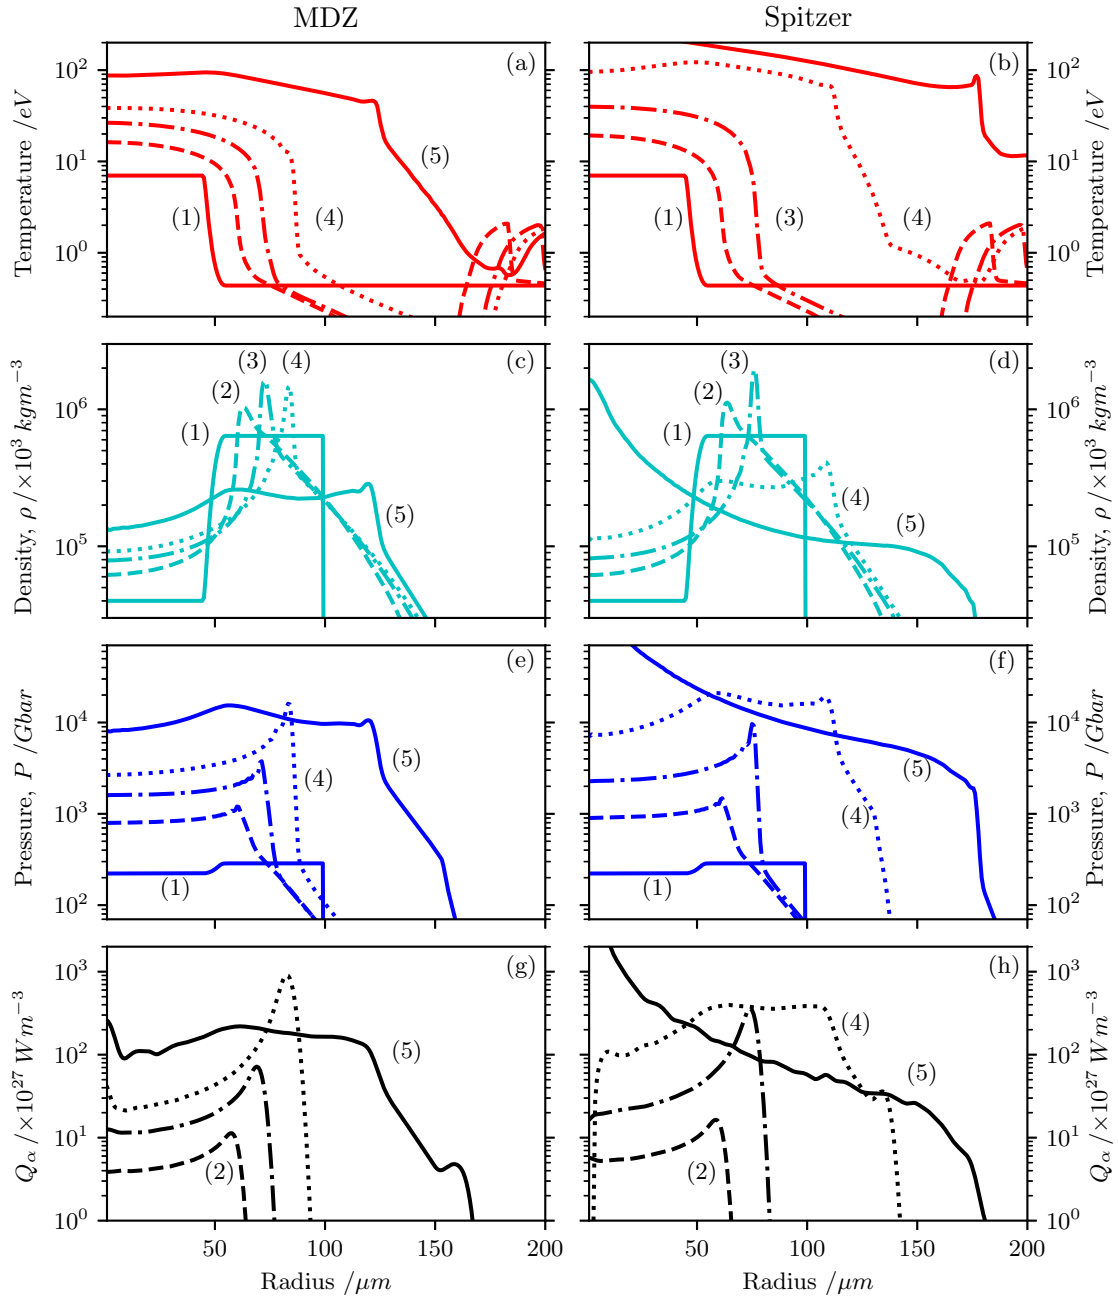
Figure 3.9: 1D radial profiles of (a,b) ion temperature, (c,d) density, (e,f) pressure and (g,h) alpha-deposition at (1) $t = 0ps$, (2) $t = 100ps$, (3) $t = 120ps$, (4) $t = 130ps$ and (5) $t = 140ps$ from a Chimera simulation of the isobaric hotspot ignition and burn, with the MDZ (left) and Spitzer (right) stopping models. Initial conditions for the hotspot are as in figure 3.8: $\rho = 4 \times 10^4 kgm^{-3}$, $T_i = 7keV$.

to the higher stopping-power of the Spitzer model resulting in higher alpha-heating deposition. As a result, the hotspot heats up faster and loses less heat to the shell, and therefore ignites and burns faster. We note that the different stopping model does not seem to affect the qualitative evolution of the hotspot from assembly through ignition to burn here, but merely accelerates (or decelerates) the process.

The rate at which the hotspot evolves under the Spitzer model is more similar to Atzeni's results, noticeable particularly in the final time profiles (5) in both pressure and density. This is primarily due to Atzeni's use of a model which is also based on the Spitzer slowing coefficients [179].

The stronger edge-like and peak-like features in both the MDZ and Spitzer profiles relative to the Atzeni profiles are likely due to the kinetic nature of our model, compared to a multi-group diffusion model. Diffusive models are not non-local, and as such struggle to capture features such as the Bragg peak in the energy-deposition profiles [91]. Although the multi-group aspect brings the modelling into better agreement with full non-local treatments (such as the Monte-Carlo model as has been implemented in Chimera), it is still less capable of capturing non-local behaviour.

4 1D Hotspot Behaviour

In this chapter we explore the behaviour of the hotspot in one dimension (i.e. radial dependence only), particularly in relation to its power balance. Section 4.1 describes the 1D simulations of the ablation and implosion phases which establish the platform for the work in this thesis. In section 4.2, the validity of the assumption that the hotspot is optically thin is investigated. Section 4.3 illustrates the various burn regimes through which a hotspot can evolve, as described in section 2.4. The properties and power balances of the hotspot in each regime are compared. In section 4.4, we quantify the error in the hotspot calculations described in section 4.3 by examining various definitions of the hotspot and comparing their properties and power balances in the various burn regimes.

4.1 Ablation Phase Calculations

The 1D simulations of the ablation and implosion phases are described in this section. Much of the work to ensure accuracy and reliability in these simulations was done by K. McGlinchey, and described in his thesis work [130]. The simulations were run *ab initio*, can be driven through either a radiation temperature pulse, or a frequency-dependent spectrum which includes non-Planckian emission properties, such as the *m*-band emission from the hohlraum [130]. The radiation transport for these simulations are run using 54 frequency groupings, chosen non-uniformly in order to capture spectral detail such as ionisation edges and line structure from the ablator in order to ensure the correct interaction between the radiation drive and the capsule, as described in [130]. Work by McGlinchey [130] includes adjusting the radiation drive to ensure the correct shock timings. Comparisons between results using Chimera and HYDRA are used [130] to ensure the accuracy of our simulations of NIF implosions, matching criteria such as implosion velocity, fuel adiabat and neutron yield.

Figure 4.1 shows a density contour map as a function of radius and time. The entire implosion is shown and the 4 stages of the implosion highlighted: ablation, implosion, stagnation and burn. The positions of the material interface between the DT and the ablator is shown alongside the first and second shocks.

The radiation drive incident on the ablator (from the top of the figure) drives the implosion, launching the shocks inwards. The shock merger of the 1st and 2nd shocks at the edge of the DT ice marks the implosion phase, in which the shell accelerates inwards and compresses the inner DT gas. The hotspot formation denotes the beginning of the stagnation phase, in which the hotspot decelerates the shell. During this process, the
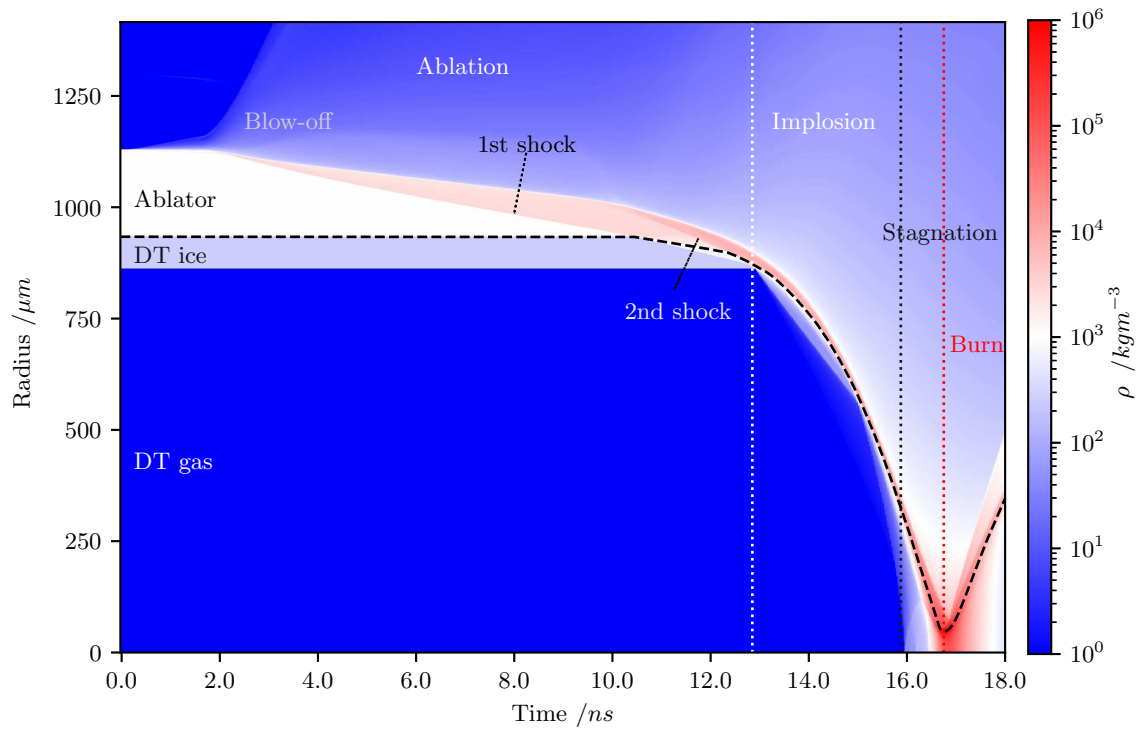
Figure 4.1: A contour map of density as a function of radius and time, showing the entire capsule implosion broken down into 4 stages: ablation, implosion, stagnation and burn. The first and second shocks propagating through the capsule are identified, and the black dashed line indicates the material interface between DT and the ablator.

shell does work on the hotspot, further compressing and heating it until the shell is completely decelerated (or stagnated). This marks the beginning of the burn phase, in which an igniting hotspot heated by sufficient alpha-heating would begin to burn the surrounding shell.

In these original simulations, no alpha-heating is included, and therefore no burn phase is included either. The focus is on ensuring the correct ablation characteristics and implosion behaviour for a given implosion set-up (in this case, N130927). The work in this thesis concentrates on the hotspot behaviour during the stagnation and burn phases of the implosion.

4.2 Radiative Losses

In this section, we explore the radiative losses in the hotspot. The hotspot is often modelled as being optically-thin (e.g. in [180]), owing to the relatively high temperature and low density (compared to the shell). Under this assumption, the hotspot bremsstrahlung free-streams out of the hotspot, with all the emitted radiative power leaving the hotspot. Such an assumption allows for both theoretical and computational simplification: theoretically, the radiative losses from the hotspot can be calculated simply from the bremsstrahlung emissivity; and computationally, the radiation transport model is computationally-intensive, requiring large numbers of subcycles ($\sim 10^2 - 10^3$) (owing to the short cell-transit time of radiation compared to the hydrodynamic timescale) for ~ 50 frequency groups. Here, we examine the validity of this assumption by comparison to a full radiation transport model, which includes contributions from emission and absorption as the radiation is transferred through the plasma.

4.2.1 Radiation models and simulation setup

We compare the results of simulations using two separate radiative loss models: one using the full radiation transport model (as described in section 3.1), and an "elementary emission" model. This elementary emission model considers radiation as purely lost from the hotspot, where the type of radiative emission is dependent on the opacity. This is done first by calculating the frequency-integrated bremsstrahlung emissivity, J using equation 2.29. Then, using the black-body emissivity:

$$U_p c = \sigma T^4 \tag{4.1}$$

where U_p is the Planckian radiative energy density, we calculate a frequency-averaged opacity $\kappa = J/cU_p$. This is used to calculate the optical depth of a cell of size Δx:

$$\tau = \kappa \Delta x \tag{4.2}$$

This optical depth is used to determine whether a given cell emits as an optically-thick black-body surface (i.e. using equation 4.1) or as an optically-thin bremsstrahlung
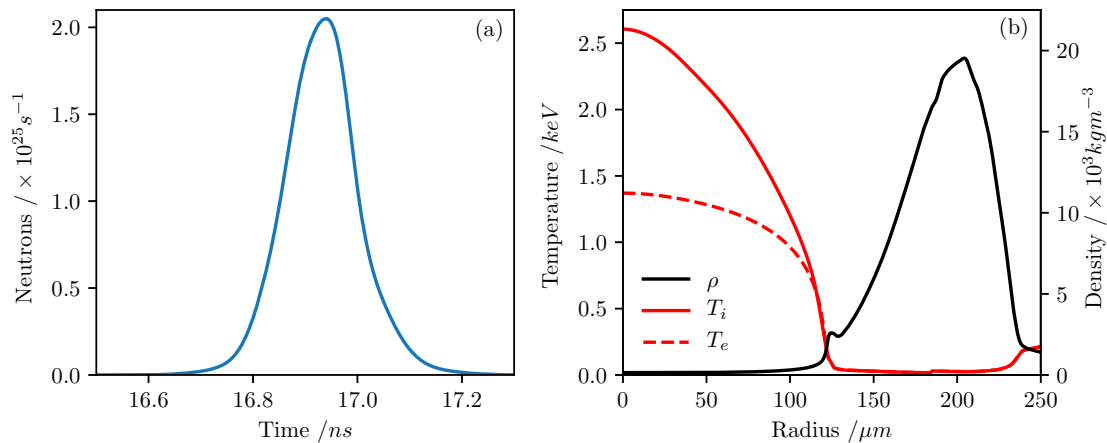
Figure 4.2: (a) The neutron burn history from the original 1D data set, and (b) radial profiles of density (black), ion temperature (red, solid) and electron temperature (red, dashed) at $t = 16.475ns$.

volume (i.e. using equation 2.29). This elementary emission model does not incorporate any method for absorption of radiation, only emission. The full-radiation transport model is as described in 3.1.

The simulations in this section are based on the NIF shot N130927, and are restarted from a 1D dataset generated as described in the previous section. Both simulations are restarted from the same data at the same point in time of 16.475ns, just before the start of the neutron pulse. The restart conditions are shown in figure 4.2, showing (a) the neutron burn history and (b) the radial density and temperature profiles at the time of restart.

4.2.2 Model comparison

The elementary emission model resulted in a DT neutron yield of 8.45×10^{14}, while the full transport model gave 3.47×10^{15}. This difference in yield is significant, with the radiation-transport yield around 4× that of the elementary emission model.

We explore the differences in the implosions in figure 4.3, showing the contour maps as a function of time and radius for the density and ion temperature for both models. Considering the density maps, the dense shell (in red) for the full transport decelerates at around 16.7ns and rebounds off the forming hotspot at around 16.9ns. In contrast, the shell in the elementary emission model continues to implode towards the origin, converging on the origin at just past 17.0ns.

The radiation emission from the hotspot in the elementary emission model removes energy density from the hotspot, reducing the temperature and the pressure. The degree of this radiative loss is so significant that instead of decelerating the shell and heating up, the hotspot continues to cool and be compressed by the shell, becoming indistinguishable from the shell in the process. This is a stark difference compared to experimental results, in which hotspot formation is clearly observed. One can conclude
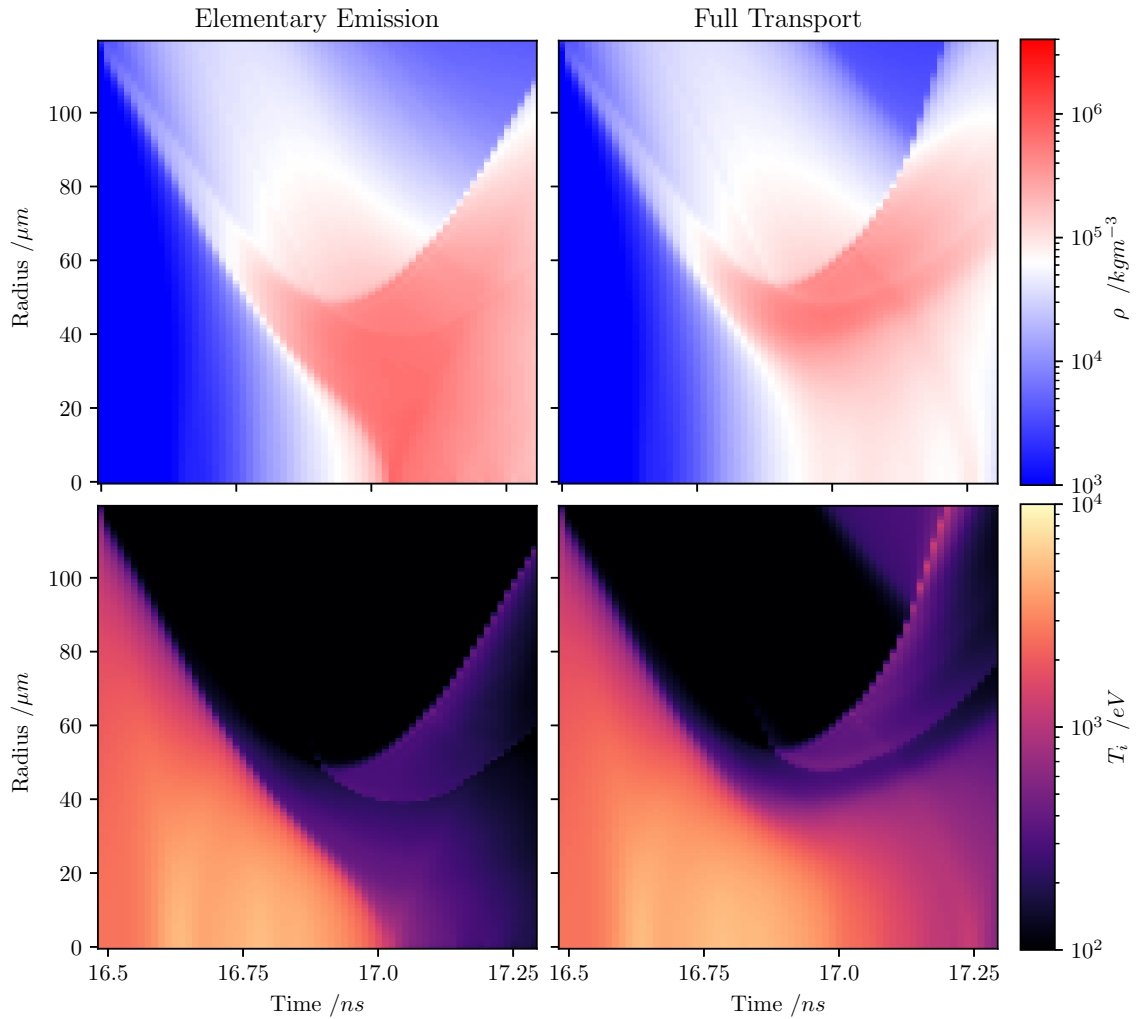
Figure 4.3: Density (top) and ion temperature (bottom) contour maps over time and radius for the elementary emission model (left) and full radiation transport model (right). The maps are shown on the same logarithmic colour scales for ease of comparison and clarity of detail.
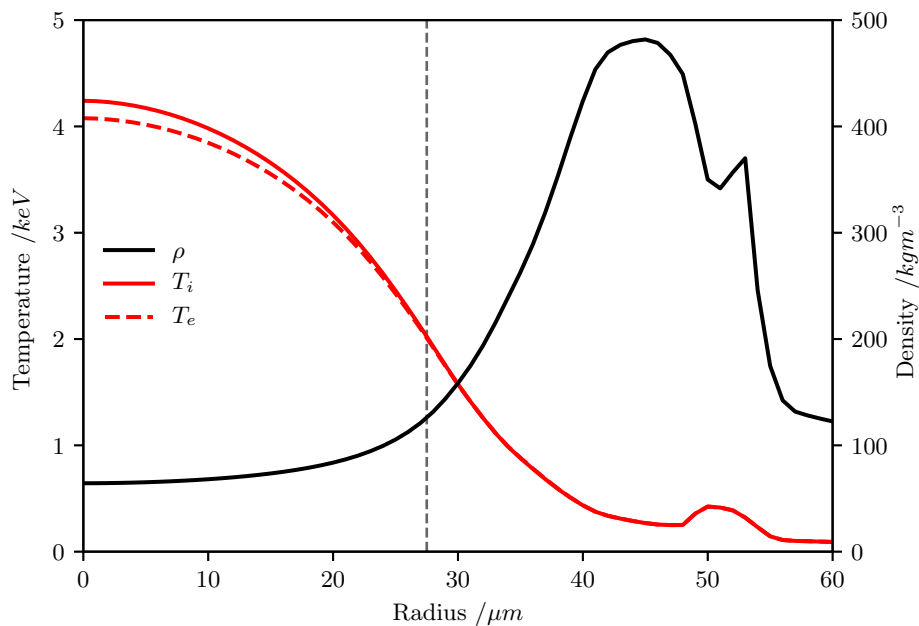
Figure 4.4: Radial profiles of the hotspot at the time of peak neutron output, bang time (16.925ns), showing the density (black), ion temperature (red, solid) and electron temperature (red, dashed). The dashed line marks the hotspot radius, as defined by the 2keV ion temperature contour.

that the hotspot is therefore not entirely optically-thin, and must be reabsorbing some of the bremsstrahlung emission, as the elementary emission model causes too high a radiation loss.

The nature of the radiation transport in the hotspot is complex — each hotspot cell has a certain emissivity depending on its density and temperature, but also reabsorbs some of this self-emission as well as incoming radiation from surrounding regions, with the opacity also dependent on density and temperature. Rather than examining the opacity and emissivity profiles of the hotspot, we consider the simplified but illustrative problem with only the radiation emitted from the centre of the hotspot at bang time (i.e. for a single time snapshot — the full problem is also time-dependent), where bang time is defined as the time of peak neutron output.

Figure 4.4 shows the radial density and temperature profiles at bang time for the simulation. Figure 4.5(a) shows the spectrum of the source function from the centre of the hotspot — effectively the emission spectrum from the cell — compared to the optically-thin bremsstrahlung emission spectrum ($\propto \exp(-h\nu/k_bT)$ [181]) and the optically-thick black-body spectrum. The peak emission in the spectrum is at $\sim 670eV$, and is a result of the plasma at the centre of the hotspot being optically-thick to low-energy radiation. For low photon energies, the plasma is optically-thick and emits as a black-body, but the spectrum transitions to a bremsstrahlung emission spectrum as the photon energy increases.

Figure 4.5(b) shows the cumulative optical depth from the centre of the hotspot as a function of radius and photon energy. This is effectively the cumulative spatial integral
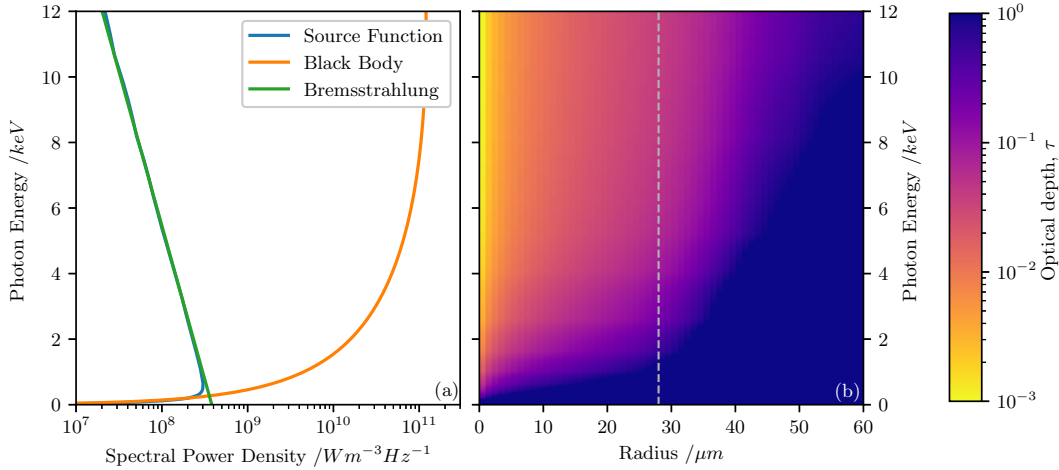
Figure 4.5:   (a) The spectrum of the source function (blue) in the central cell of the hotspot, compared to the optically-thin bremsstrahlung emissivity (green) and the optically-thick black-body (orange) and (b) the integrated optical depth as a function of photon energy and radius. This is a running integral of opacity over radius from the centre of the hotspot to the radial value. The dashed line indicates the hotspot radius, R_{hs}.

of the opacity, i.e.

$$\tau_{i,\nu} = \int_0^{x_i} \kappa_\nu(x) dx \tag{4.3}$$

Beer's law describes the intensity of radiation of a given energy, I_ν at a given radius:

$$I_\nu(x = x_i) \;=\; I_\nu(x = 0) e^{-\tau_{i,\nu}} \tag{4.4}$$

The limits of the colour bar range from $\tau = 10^{-3}$ (i.e. effectively optically-thin, with little to no absorption of this photon energy at this point) to $\tau = 1$. At $\tau = 1$, 63% of the radiation from the centre of the hotspot at this photon energy has been absorbed by the time the radiation has travelled outward to this radius.

For radiation at $1keV$, this radius is about $25\mu m$, and around $40\mu m$ for $4keV$ photons. Comparing the spectrum from figure 4.5(a), 18% of the radiative energy emitted from the centre of the hotspot is at photon energies $< 1keV$ which is reabsorbed before the edge of the hotspot is reached.

From this simplified scenario, we can then consider figure 4.6, which shows the radial profiles of the net radiation loss from each of the elementary emission and the full radiation transport models at bang time, in addition to the emission and absorption contributions to the full transport model. Note that the full transport emission and the elementary emission profiles are similar in the centre of the hotspot, up to around $20\mu m$. This is to be expected, as the centre of the hotspot is where the optically-thin assumption is most valid. However, the absorption in the centre is not insignificant, although still small ($\sim \frac{1}{3}\times$ the emission). The contribution from the absorption increases significantly towards the edge of the hotspot, as does the emission. This is due to the higher density
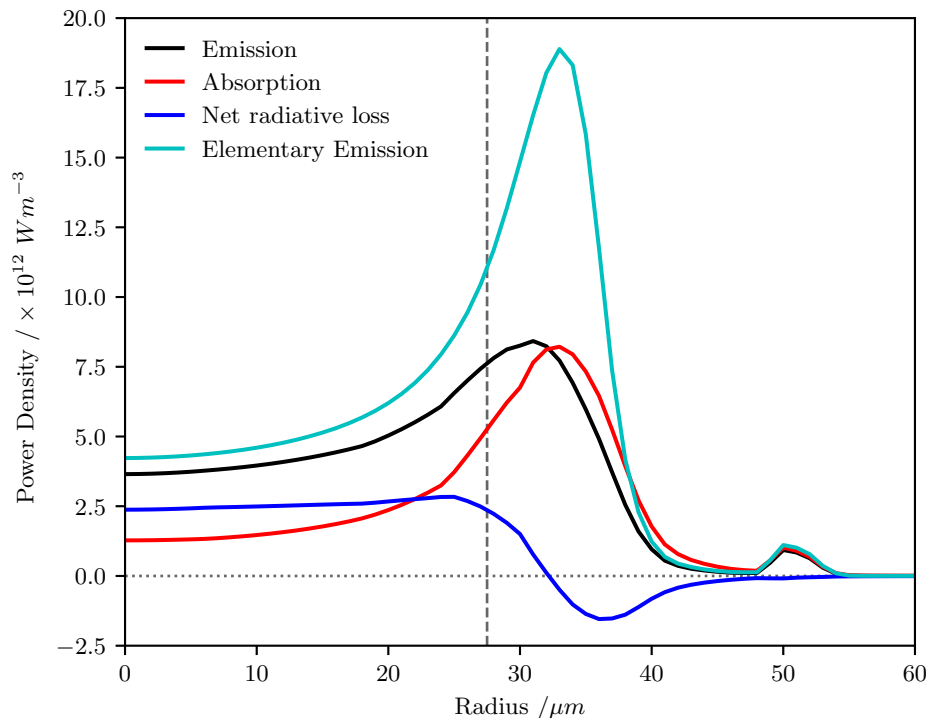
Figure 4.6: Radial profiles at bang time of the net radiation loss from the bremsstrahlung model (cyan) and the full radiation transport (blue), showing the frequency-integrated emission (black) and absorption (red) contributions.

— both emission and absorption scale as n^2. The emission then drops further out as the temperature drops, with the absorption following suit.

The net radiative loss (emission − absorption) is significantly lower than the optically-thin model in the centre ($\sim \frac{1}{2}$). Even as the emission increases towards the edge of the hotspot, the corresponding increase in absorption results in the net radiative loss remaining at a relatively constant power density throughout most of the hotspot. As the emission falls off beyond the edge of the hotspot, the absorption overtakes it and results in a net absorption of radiation in the denser regions of the shell.

Overall the full radiation transport profile differs significantly from that of the optically-thin model, particularly at the edge of the hotspot. Although the elementary emission model could, in theory, be artificially reduced in order to produce better agreement with the emission from the full transport, it would still be unable to reproduce the full radiative behaviour from the absorption of radiation. The fundamental inability to account for reabsorption of radiation results in an overall volume-integrated power loss that is $3\times$ higher than that of the full transport model. Under the elementary emission model, the volume-integrated hotspot power loss is equal to $1.08W$, while the full transport model gives a much lower $0.38W$. This is broken down into $0.83W$ emitted and $0.45W$ reabsorbed within the hotspot, such that the overall emission is less than half of the that emitted. Clearly, under the full transport model the hotspot reabsorbs a significant proportion of its own emission. In addition, the hotspot-shell boundary region is opaque

enough to result in significant absorption of radiation. The overall profiles indicate net transport of energy from the hotspot to the shell, and must therefore be considered as this can result in mass ablation, increase density- and temperature-gradient scale lengths and reduce instability growth.

4.3 Hotspot Burn Regimes

We now explore the properties of the hotspot as it progresses through the various regimes of burn as described in section 2.4. Simulations in this section are again based on NIF implosion N130927, starting from a $0.25\mu m$ resolution 1D dataset run with 54 non-uniform radiation groups and a frequency-dependent x-ray drive spectrum which includes the m-band component of the radiation flux [8]. All simulations are restarted at $0.5\mu m$ resolution from the same dataset at the same time (the time of peak radiation temperature, 15.7ns), and hydrodynamic scaling (section 2.5) is used to access the various regimes while keeping the implosion design similar. This scaling is applied at the time of restart, with the initialisation data (temperatures, densities and velocity) effectively radially stretched in order to match the scaling requirements, while the magnitudes remain unchanged. These simulations are run with the MDZ stopping model in the alpha-particle transport, and explicit thermal conduction.

Recall the descriptions of the various burn regimes from section 2.4: in the self-heating phase, there is only a low level of alpha-heating present, which primarily deposits heat within the hotspot, but has little impact overall on the hotspot temperature or evolution. In 'robust ignition', alpha-heating boosts the temperature significantly, thereby increasing the fusion reactivity and causing yield amplification. However, the confinement is insufficient to allow burn to develop. In propagating burn, the hotspot ignites robustly with adequate ρR to confine it, and as a result a burn wave develops and propagates outward through the confining shell. The gain in yield occurs due to the ablation of mass into the hotspot, corresponding to an overall increase in burning material.

Figure 4.7 shows the time-evolution of radial profiles of a scaled down, $S = 0.9$ N130927 simulation. Profiles are shown for the fuel (DT) and ablator (CH) densities, ion and electron temperatures, and the contributions to the hotspot power balance from alpha-heating (Q_α), thermal conduction (Q_e), radiation (Q_γ) and mechanical work (Q_m). The vertical dotted line indicates the hotspot radius, R_{hs}, as defined by $T_i \geq 2keV$. Note that the next section will explore different potential hotspot definitions and their differences. Times are shown relative to the time of peak compression, t_0. Figures 4.8 and 4.9 show the time-evolution of the radial profiles for $S = 1.0$ and $S = 1.1$ capsules exhibiting the robust ignition and propagating burn regimes respectively. The timesteps between panels are equal when scaling is accounted for, and the radial scale is also adjusted by the scale factor S. These allow for direct comparison across the figures in order to compare behaviour between regimes.

In the self-heating regime, from figure 4.7, we can see that after peak compression, the alpha-heating contribution to the power balance is less than the combination of the
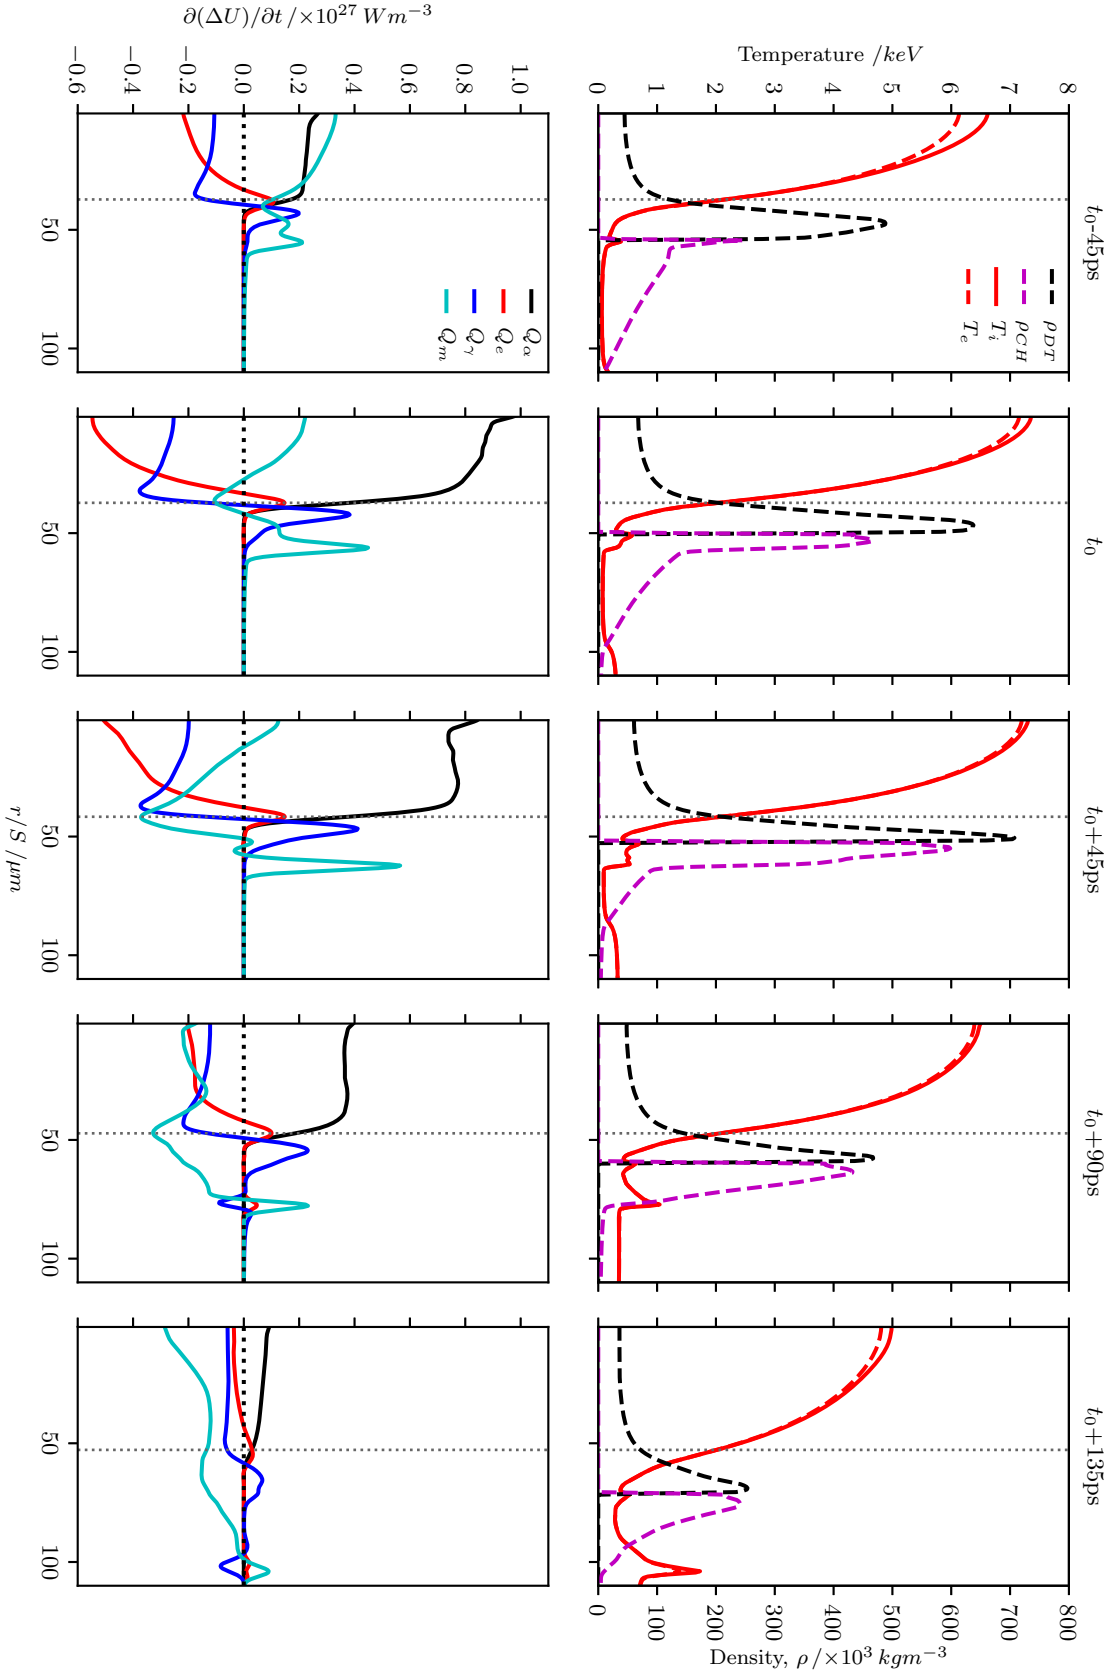
Figure 4.7: Time evolution of a 1D implosion in the self-heating regime based on shot N130927, at scale factor 0.9. Times are relative to peak compression at t_0. Densities and temperatures (above) are shown as fuel (dashed black) and ablator (dashed cyan) densities, ion temperature (solid red) and electron temperature (dashed red), while volumetric contributions to the hotspot power balance (below) are broken into alpha-heating, Q_α (black), electron thermal conduction, Q_e (red), radiation, Q_γ (blue) and mechanical work, Q_{P_dV} (cyan). The zero value is shown in the dashed black line, and the hotspot radius, R_{hs} is given by the vertical dotted line.
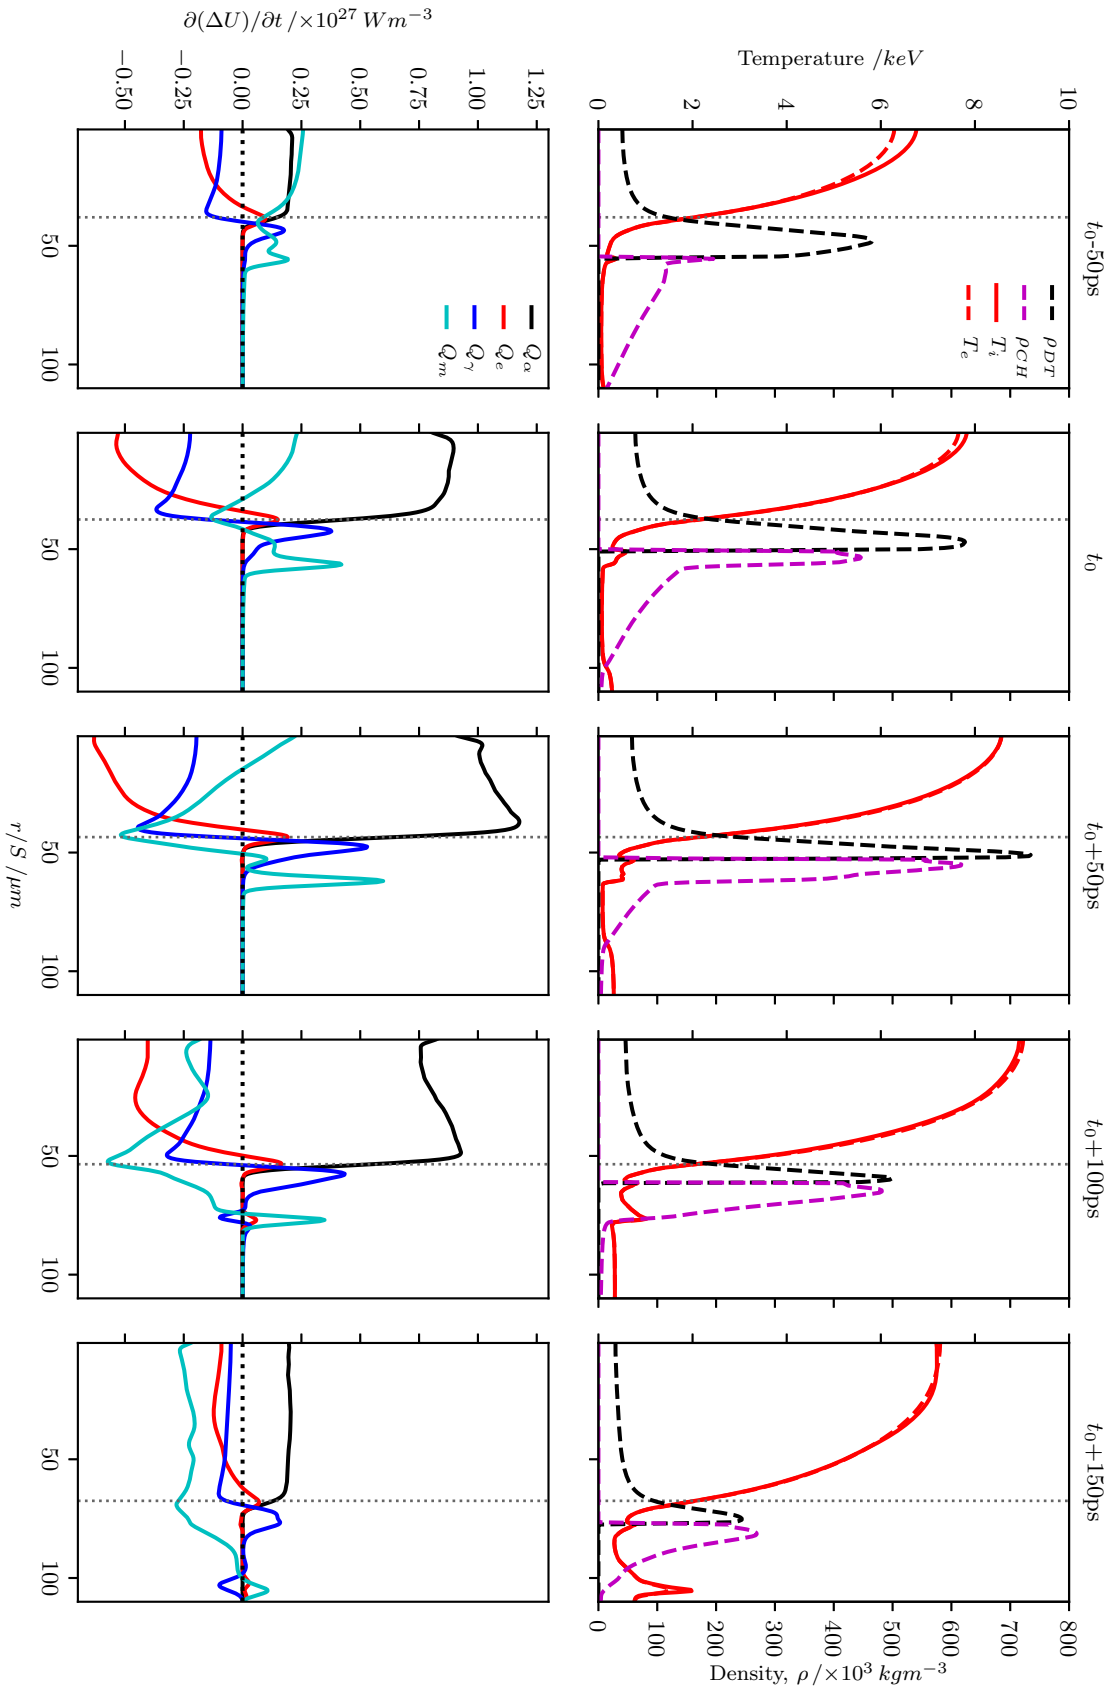
Figure 4.8: Time evolution of a 1D implosion undergoing robust ignition, based on N130927 at scale factor 1.0 as shot. Quantities are as in figure 4.7. Time-steps are proportional to scale factor to allow direct comparison between figures 4.7, 4.8 and 4.9.
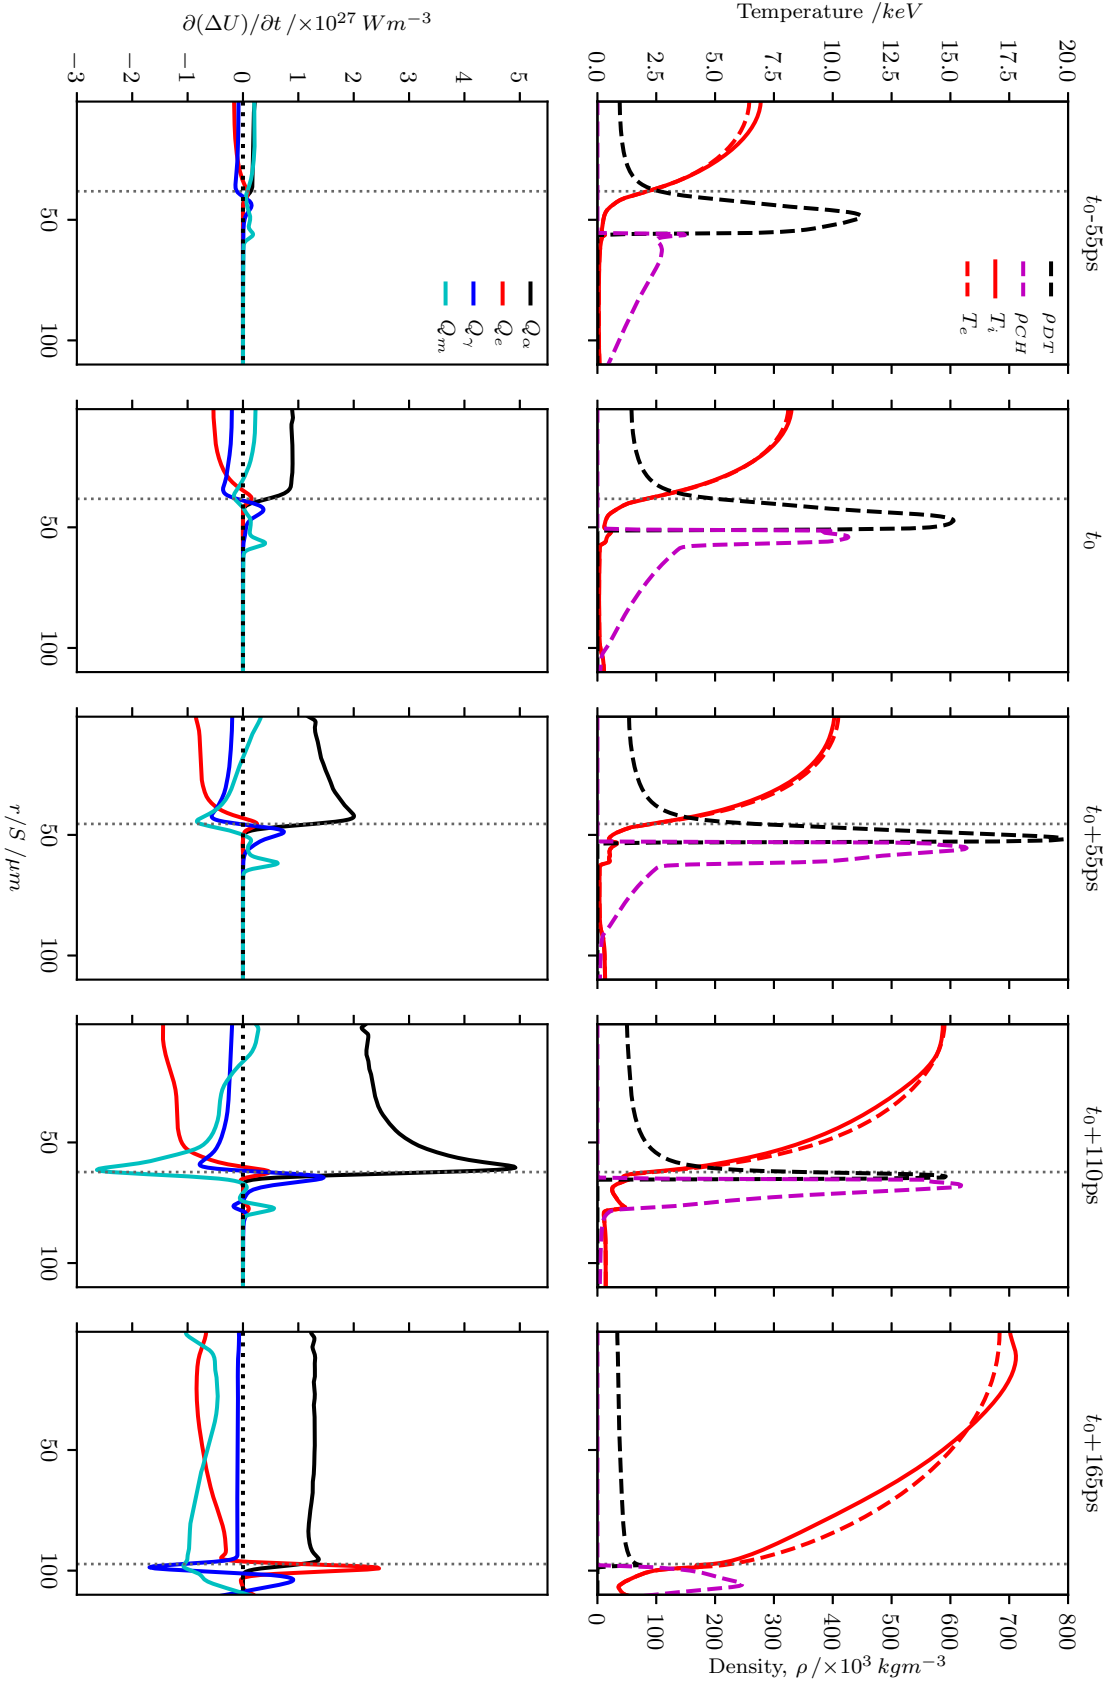
Figure 4.9: Time evolution of a 1D implosion in the propagating burn regime based on shot N130927, at scale factor 1.1. Quantities are as in figure 4.7. Time-steps are proportional to scale factor to allow direct comparison between figures 4.7, 4.8 and 4.9.

losses from thermal conduction, radiation and mechanical expansion. As a result of this power balance, the hotspot temperature does not increase after peak compression, but gradually falls. The mean free path of the alpha-particles, l_α (i.e. the alpha-range) is less than the size of the hotspot, such that $\approx 80 - 85\%$ of the alpha-heating energy is deposited within the hotspot. Defining the transparency to alpha-particles, τ_α [126, 127]:

$$\tau_\alpha = \frac{l_\alpha}{R_{hs}} \tag{4.5}$$

for a hotspot of radius R_{hs}, the transparency of the self-heating hotspot is relatively constant throughout the burn pulse at around 75%.

In this regime, the total hotspot heating timescale is long relative to the confinement time of the hotspot. Even though the majority of the alpha-heating is deposited in the hotspot, there is not enough alpha-heating overall to raise temperatures significantly or to overcome the losses by other mechanisms. There is little mass ablation of the shell (alpha-driven or otherwise), and the overall fraction of fuel mass contained within the hotspot remains low, at $< 20\%$.

Figure 4.8 illustrates the time-evolution of a capsule in the robust ignition regime, using an at-scale ($S = 1.0$) simulation of N130927. Here, stronger alpha-heating levels boost the hotspot temperature significantly, with the central temperature $1 keV$ hotter than for the self-heating regime at peak compression, and continuing to increase after peak compression. The mechanical work produces a hotspot of a similar temperature from hydrodynamic scaling arguments, but the larger size (and therefore areal density) of the hotspot results in better retention of any initial alpha-heating produced. In addition, the larger scale of the capsule results in greater levels of fusion, and therefore more alpha-heating. The resultant alpha-heating is enough to increase the temperatures by around $1 keV$ during the $100 ps$ after peak compression, even against the hotspot expansion losses. However, the alpha-heating is not enough to result in significant heating beyond the hotspot — the areal density of the shell is not enough to confine this robustly igniting hotspot. The resultant increase in temperature increases the fusion reactivity and leads to a resultant boost in yield (from 4×10^{16} to 1.4×10^{17}), but the lack of confinement leads to only a modest increase in hotspot mass fraction to around $\sim 30\%$ from fuel ablation. This leads to a burn-up efficiency of $\Phi \lesssim 1\%$, where we can define the burn-up efficiency [7]:

$$\Phi = \frac{N_{fus}}{N_{DT}^{(0)}} \tag{4.6}$$

N_{fus} is the total number of fusion reactions, and $N_{DT}^{(0)}$ is the initial total number of DT pairs present in the capsule.

The alpha-heating profile shifts from the self-heating profile (with energy deposition relatively constant throughout the hotspot, and sigmoid-like at the edge of the hotspot) towards a Bragg-peak like profile throughout the implosion. This is due to the increase in hotspot temperature from heating, and the decrease in density due to expansion,

both of which result in a larger alpha-range. In addition, more alpha-particles are being produced at larger radii, i.e. closer to the shell, and therefore overall more alpha-particles will be reaching and thermalising in the dense fuel layer. The proportion of alpha-heating energy deposited in the hotspot decreases from $\approx 80\%$ to $\approx 75\%$. The increased ablation and hotspot pressure result in sharper density- and temperature-gradients.

The evolution of the propagating burn regime illustrated by figure 4.9 shows the propagation of a burn wave into the dense shell. The temperature in the hotspot continues to increase even as the capsule expands, building the pressure in the hotspot and resulting in a faster re-expansion rate. The alpha-heating contribution is significantly greater than the combination of the loss terms (from radiation, thermal conduction and mechanical expansion). The profile evolution of Q_α from heating primarily the hotspot (as in the self-heating regime) through a small Bragg peak (as in robust ignition) to a strongly Bragg-peaked profile can be seen in the middle 3 timesteps of figure 4.9. This can be explained as an extension of the earlier Bragg peak formation in the robust ignition profiles; with stronger fusion production in close proximity to the shell (i.e. more alpha-particles being born in the outermost regions of the hotspot, rather than the predominantly in the centre in the earlier regimes), more alpha-particles experience the strong stopping power in the dense shell and deposit energy abruptly in the thin layer. As a result, the alpha-heating (in conjunction with an increased thermal conduction) propagates the burn wave, driving significant mass ablation. Note that the radiative deposition into the fuel ahead of the alpha-heating and thermal conduction heat flows effectively pre-heats the fuel for the deflagration wave. This allows the alpha-heating to penetrate deeper into the shell, which could act to increase the propagation speed of the burn wave. The ρR of the shell is high enough to adequately confine the high-pressure burning hotspot, thus compared to robust ignition, where the yield amplification is due to the increase in hotspot temperature, the hotspot propagates a burn wave into the shell, ablating material and increasing the hotspot mass, and therefore the mass of burning material. The hotspot mass fraction increases to include almost the entirety of the fuel for this $S = 1.1$ simulation, leading to a burn-up of $\Phi \approx 5\%$ (corresponding to a DT neutron yield of 1.7×10^{18}). The temperature- and density-gradients are sharper than the robust ignition case, again due to even more ablation and higher pressures.

As the alpha-heating contribution to the power balance increases from the self-heating regime to the propagating burn regime, the thermal conduction and radiation contributions also increase, scaling as $\sim T^{7/2}$ and $\sim T^{1/2}$ (for a simple bremsstrahlung emission) respectively (see section 2.2). The stronger alpha-heating regime and respective temperature scalings result in the increased importance of thermal conduction relative to the radiation loss. The thermal conduction transports heat from the centre of the hotspot, where it is hottest, to the edges, where it is deposited into the closest region of the shell. The radiation loss instead peaks near the outer edges of the hotspot, where the density begins to increase and before the temperature has become too cold, and is deposited further out into the shell than the thermal conduction. A significant proportion of the

overall radiation is reabsorbed in the dense shell, reducing the amount that escapes the capsule completely. Since much of the energy from the hotspot loss in these processes heats up shell material, these losses could be considered as recycled [111] if the corresponding material is ablated into the hotspot. This is because mass ablated into the hotspot must be heated up, and therefore the heating of the shell material from these losses reduces the heating required when this material is ablated into the hotspot. The thermal conduction is deposited just outside the hotspot, and thus the thermal heat flow is easily recycled by a hotspot propagating a weak burn wave. Recycling the radiation heat flow would require a much stronger burn wave than for the thermal conduction case, due to the deposition being further out.

Figure 4.10 compares the neutron burn history and various hotspot properties against scaled time for the three regimes. These properties include: the fraction of fuel mass contained within the hotspot, M_{hs}/M_{DT}; the burn-averaged density, $\langle \rho \rangle$; the burn-averaged ion temperature, $\langle T_i \rangle$; the scaled areal density, $(\rho R)_{hs}/S$; and the scaled hotspot radius, R_{hs}/S. Here, a burn-averaged quantity, $\langle Q \rangle$ is the fusion-weighted volume average of Q, calculated as:

$$\langle Q \rangle = \frac{\int_{V_{hs}} QF dV}{\int_{V_{hs}} F dV} \tag{4.7}$$

where $F = \frac{1}{4} n^2 \langle \sigma \nu \rangle$ is the fusion production.

The hotspot mass fraction clearly illustrates the different regimes, showing very little mass ablation for self-heating, modest ablation for robust ignition and significant ablation for propagating burn. The propagating burn mass reaches almost 100% of the DT fuel in this example of the regime.

The burn-averaged densities and temperatures follow similar scaled time evolutions before stagnation. The stronger alpha-heating regimes are already slightly hotter by peak compression, and therefore less compressible, resulting in slightly lower densities. After stagnation, the self-heating temperature gradually begins to fall, while the robust ignition temperature continues to increase very slightly. The propagating burn regime results in much starker differences in the time-evolution. The density reaches a similar peak as for self-heating and robust ignition, but then is maintained at a similar level due to ablation, even as the others fall due to re-expansion. Here we can see that the hotspot density does not actually appear to increase from the mass ablation, as the expansion due to hydrodynamics (and to a lesser extent, the burn propagation) produces an opposing effect. Instead, the time-width of the peak in density is broader. The temperature continues to build as more and more fusion occurs, demonstrating the central ignition bootstrap process.

The hotspot radii are initially exactly the same (when adjusted for the scale factor), as expected from hydrodynamic scaling. However, they diverge as the differing levels of alpha-heating begin to contribute to heating the hotspot, increasing the pressure and accelerating expansion. The robust ignition expansion is only mildly faster than that of the self-heating regime, but significantly faster for propagating burn. This is
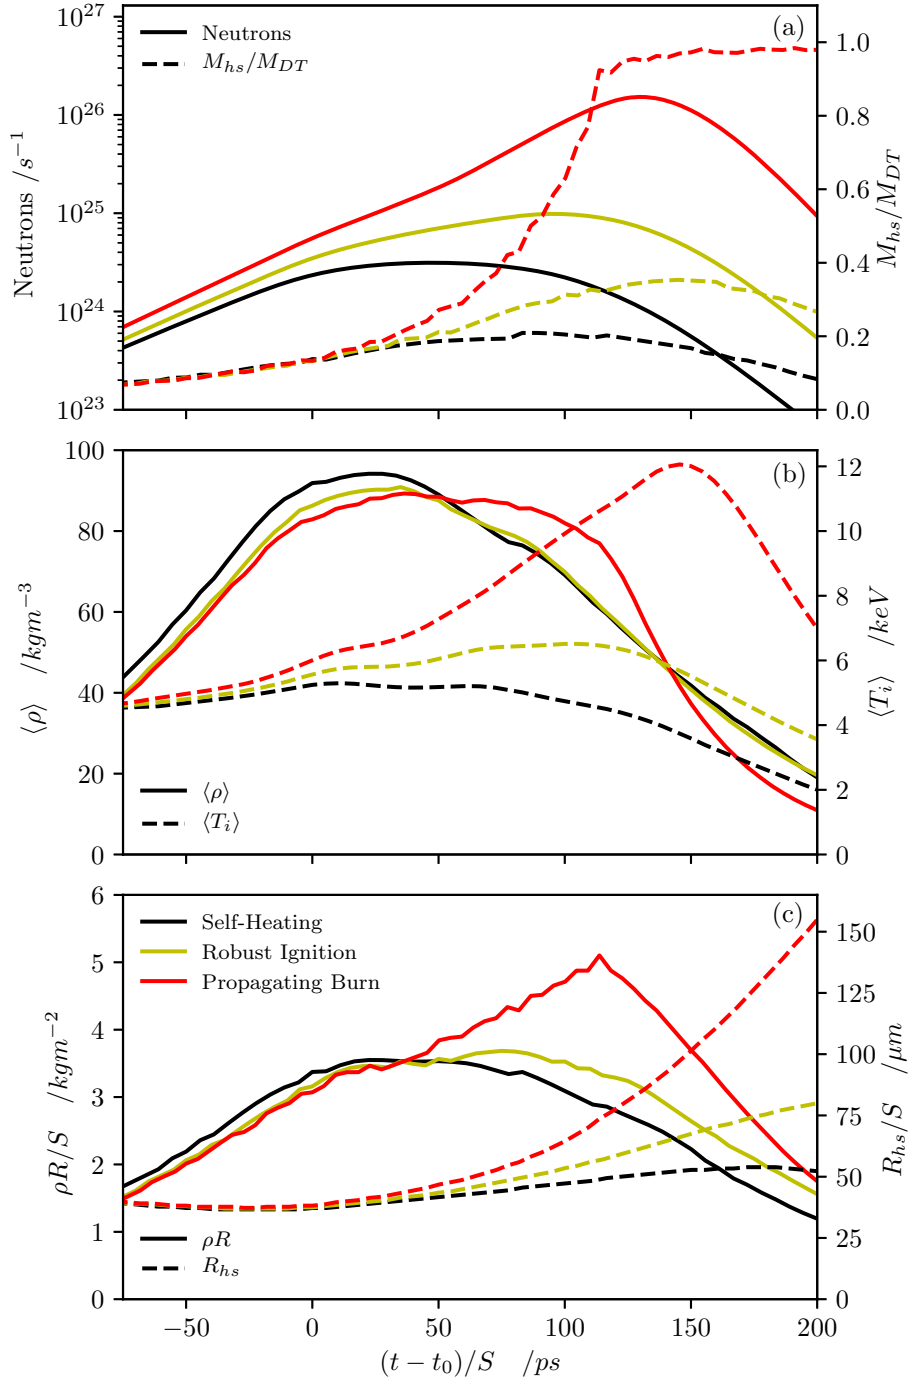
Figure 4.10: (a) Neutron burn history (solid) and fraction of DT mass in the hotspot (M_{hs}/M_{DT}) (dashed); (b) hotspot burn-averaged density, $\langle \rho \rangle$ (solid) and ion temperature, $\langle T_i \rangle$ (dashed); and (c) scaled hotspot areal density, $(\rho R)_{hs}/S$ and radius R_{hs}/S against scaled time, shown for the three regimes of alpha-heating: self-heating (black), robust ignition (yellow) and propagating burn (red).
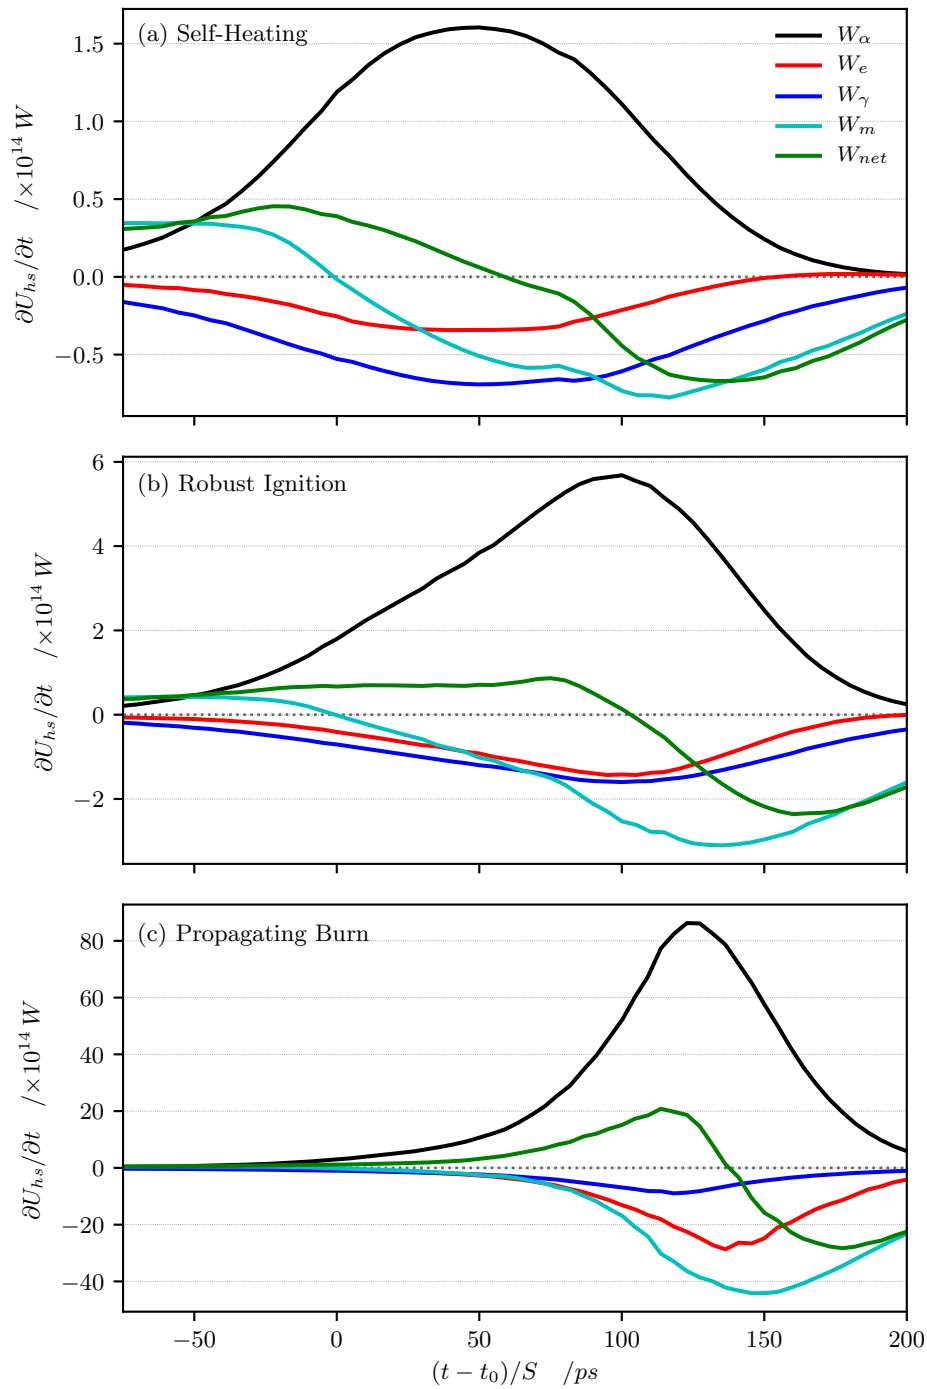
Figure 4.11: The volume-integrated hotspot power balance, $\partial U_{hs}/\partial t$ for the simulations in the (a) self-heating, (b) robust ignition, and (c) propagating burn regimes. The total (W_{net}, green) is broken into contributions from alpha-heating (W_α, black), thermal conduction (W_e, red), radiation (W_γ, blue) and mechanical work (W_m, cyan).

due both to the increased pressure causing faster hydrodynamic expansion, and the inherent increase in hotspot size due to burn propagation. The hotspot areal density can be explained from the variations in hotspot density and radius. The self-heating areal density increases up until stagnation, then gradually begins to fall off. Robust ignition follows a similar trend, but increases slightly with the faster hotspot expansion. The propagating burn areal density simply continues to increase past stagnation, as the stagnation hotspot density is maintained for a significant period of time ($\sim 100ps$) while the hotspot radius increases significantly.

Figure 4.11 shows the volume-integrated hotspot power balance for all three burn regimes discussed, including the individual contributions from alpha-heating, thermal conduction, radiation and mechanical work. In the self-heating regime, the alpha-heating is still significant with a magnitude around $2\times$ that of either radiation or thermal conduction, but not enough to also overcome the hotspot expansion. In this regime, the radiative losses are more significant than the thermal conduction losses, primarily due to the lower hotspot temperature. The re-expansion losses are similar in magnitude to the radiative losses. The stronger alpha-heating of the robust ignition regime results in a net positive, relatively constant $\partial U_{hs}/\partial t$ until almost $100ps$ after stagnation, even against the capsule expansion. The radiation and thermal conduction losses are very similar in magnitude in this regime, while the capsule re-expansion is initially similar in magnitude but becomes the largest loss term. For the propagating burn regime, the alpha-heating contribution continues to build for much longer, and to a much larger magnitude than in the previous two regimes due to the igniting hotspot. The similar increase in scale from robust ignition produces an alpha-heating contribution with a peak over $10\times$ greater, compared to the $4\times$ larger peak from robust ignition to self-heating, in addition to the longer confinement time. The resultant higher temperature results in the thermal conduction contribution peaking at around $3\times$ that of the radiation. Again however, the increased pressure and resulting faster expansion lead to expansion cooling becoming the most significant loss mechanism for the hotspot for the majority of the hotspot burn.

4.4 Hotspot Definitions

Although there are various ways in which one can define the hotspot, in general, these definitions refer to broadly the same region of the capsule. In the initial foundations and designs of ICF, there was a clear boundary between the hotspot and the shell. The hotspot was characterised by a lower density and higher temperature, and the shell by the opposite higher density and lower temperature. However, in reality, such a clearly defined boundary between the hotspot and the shell does not exist, particularly given the various mechanisms of heat flow from the hotspot which result in mass flow from the shell and increase the gradient scale-lengths. As a result, a variety of definitions are used throughout current ICF literature, all of which encompass the same hot, less dense core, but by definition include differing amounts of the boundary region. In this

section, we provide a non-exhaustive list of definitions for the hotspot, explore their differences and compare hotspot parameters and power balances calculated for these different definitions using the same simulation data.

4.4.1 Definitions

We begin with a list of potential ways to define the hotspot, noting that some of these are purely theoretical definitions, and are not intended to be experimentally measurable.

- By considering the hotspot as the hottest region of the capsule, we can define the hotspot using the $1keV$ contour, i.e. the region satisfying $T_i > 1keV$ [89, 91, 182]. Such a definition is intuitive, since the expectation is that the hotspot should be hotter than the shell. It is also easy to find for simulations, and corresponds well to the region of dominant fusion, given the temperature-dependence of the reactivity. The $2keV$ contour as used in the previous section is a variation on this definition. Hotspots using this definition will be referred to by the value of the temperature contour they are defined by, i.e. **1keV** hotspot or the **2keV** hotspot.

- An experimentally measurable definition is the **17% neutron** contour [183]. Here, the hotspot is enclosed by the contour at 17% of the peak intensity in the 2D image, where 17% is high enough to distinguish from the noise levels present in the imaging system. Considering a 1D sphere of radius R_0 and constant neutron emissivity, this definition encloses approximately 95% of the volume of the sphere.

- By considering the hotspot as the volume in which the majority of the fusion occurs, we can define the **fusion** hotspot as the volume containing 90% of the total fusion output at any given time (i.e. the 90% fusion volume). Note that due to the strong-temperature dependence of the reactivity this results in a similar hotspot to one defined by an ion temperature contour. Although similar to the 17% neutron contour, which encompasses 95% of the volume of a sphere of constant emission, the hotspot emission is not constant throughout the volume. The 17% contour from a simple self-similar neutron emission profile $\propto (1 - r^2/R_0^2)$ would consist of just 58% of the total volume, while 90% of the fusion volume would require a 2% contour. This simple emission profile is intended merely to illustrate the relative size of this hotspot definition.

- Since the hotspot can be considered as the central region of low density, we can use the density to define the hotspot. Christopherson *et al.* [184, 185] define the boundary of the hotspot using $\rho(R_{hs}) = \frac{1}{2}(\rho_0 + \rho_{peak})$, where ρ_0 is the central hotspot density and ρ_{peak} is the peak shell density. Here, we shall refer to this as the **density-average** hotspot. Other possible density-related definitions could include $\rho(R_{hs}) = e \times \rho_0$ and $\rho(R_{hs}) = e^{-1} \times \rho_{peak}$. This latter definition will be referred to as the **density e-fold**.
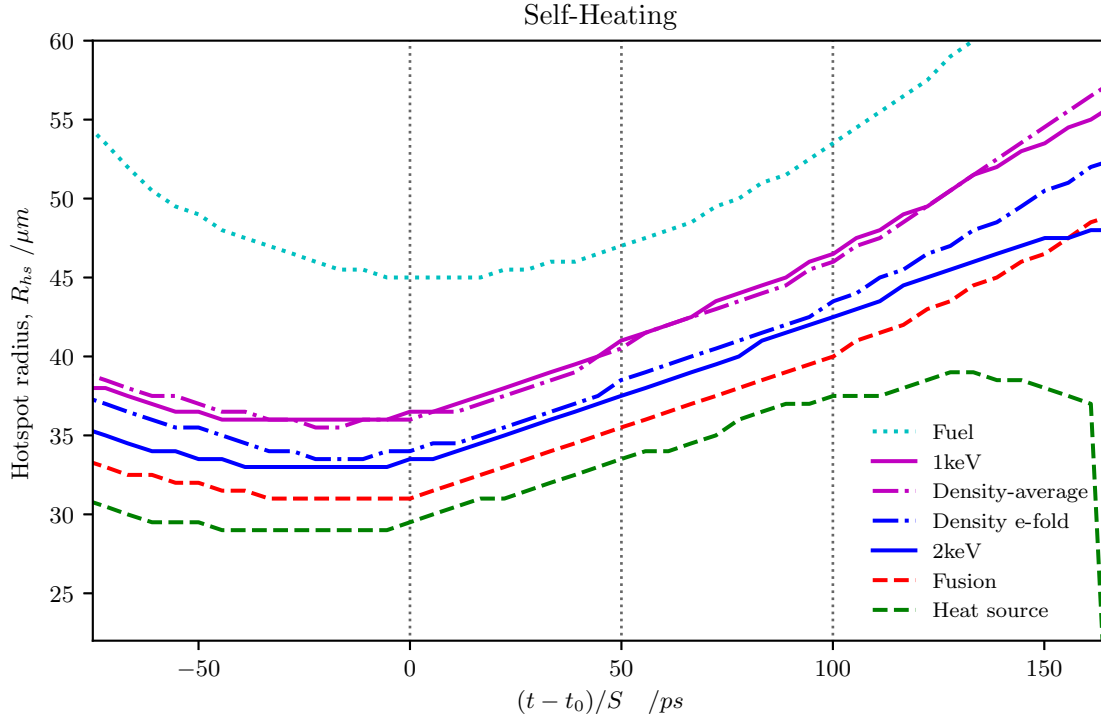
Figure 4.12: The hotspot radii against time for the scale $S = 0.9$ simulation of N130927 in the self-heating regime, for various definitions: heat source (green dashed), 90% fusion volume (red dashed), 2keV (blue solid), 1keV (magenta solid), the density e-fold ($e^{-1}\rho_{peak}$) (blue dash-dot), and the density-average ($\frac{1}{2}(\rho_0 + \rho_{peak})$) (magenta dash-dot). The edge of the fuel is also shown (cyan dotted). The time snapshots shown in figure 4.13 are indicated with the dotted lines.

- Since we expect the hotspot to heat up the rest of the shell, the **heat source** hotspot could be defined as the volume which acts as the heat source in the centre of the hotspot. In other words, this is the region, V, with the largest outward heat flow, $\int_V \nabla \cdot (\kappa \nabla T) dV$, which acts to heat the surrounding volume.

4.4.2 Comparison

We look to determine how different hotspot definitions might affect calculations of parameters of the hotspot, such as the temperature, density, radius. In addition, given the shape of the radial profiles of various contributions to the hotspot power balance, we expect that the definition of the hotspot might affect the significance of different contributions. Given the variation in hotspot behaviour with the burn regime, we also examine how the respective definitions vary.

Figure 4.12 shows the hotspot radii against time for the self-heating $S = 0.9$ simulation of N130927, for the $1keV$, density-average, density e-fold, $2keV$, fusion and heat-source hotspots, as well as the overall fuel volume. Figure 4.13 shows the radial density, temperature and power balance contribution profiles over several time step around peak compression, focussing on the edge of the hotspot. Figure 4.13 is effectively a zoomed in version of the central three time-steps of figure 4.7. The boundaries
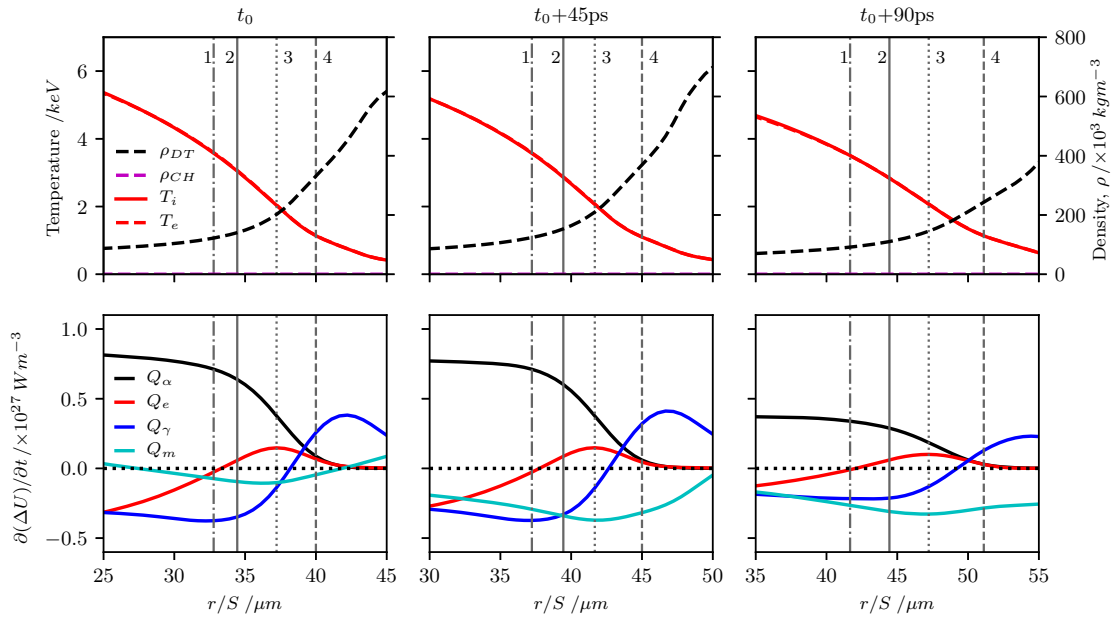
Figure 4.13: A zoomed in version of figure 4.7 in the self-heating regime, focussing on the boundary region between the hotspot and the shell. Radial profiles are shown for: (top) DT density (black dashed), ablator density (cyan dashed), ion temperature (red solid) and electron temperature (red dashed); and (bottom) the alpha-heating (black), thermal conduction (red), radiation (blue) and mechanical work (cyan) contributions to hotspot power balance. The vertical lines indicate the boundary of the hotspot as defined by: (1) heat-source (dash-dot), (2) fusion (solid), (3) $2keV$ (dotted) and (4) density-average (dashed).

Figure 4.14: Volume-integrated contributions to the hotspot power balance for the self-heating $S = 0.9$ scale simulation of N130927, showing alpha-heating (black), thermal conduction (red), radiation (blue) and mechanical work (cyan), for hotspots definitions: (a) heat-source, (b) 90% fusion, (c) 2keV and (d) density-average.

of the hotspot as defined by the heat-source, the 90% fusion volume, the $2keV$ contour and the density-average are shown. Only a few of the possible hotspot definitions are shown on this plot to retain clarity.

We see that the heat-source hotspot is the most central, since energy deposition by the thermal heat flow (Q_e) occurs at a small radius (compared to other hotspot boundaries) and high temperature ($\sim 3.5keV$). Note that beyond $\sim t_0 + 150ps$, the heat source hotspot falters as a valid definition — however, this is after the majority of the neutrons have already been produced, i.e. beyond the period of significant fusion production. The fusion hotspot is the second smallest of those shown, roughly $2 - 3\mu m$ larger than the heat source hotspot. The radii of the $2keV$ hotspot and the density e-fold hotspot follow very similar trajectories for the majority of the implosion, around $5\mu m$ larger than the heat source hotspot, as do the $1keV$ hotspot and the density-average hotspot, which are roughly $8 - 10\mu m$ larger.

In this regime, we can see that the relationships between various hotspot definitions are generally well-behaved and vary little for the most part. The boundaries of the definitions shown in figure 4.13 do not vary in their size-ranking, with the gap between boundaries varying only by a constant factor.

Figure 4.14 shows a comparison of the volume integrated contributions to the hotspot power balance from each of alpha-heating, thermal conduction, radiation and mechanical

Figure 4.15: Burn-averaged (a) ion temperature, $\langle T_i \rangle$ and (b) density, $\langle \rho \rangle$ for the self-heating $S = 0.9$ scale simulation of N130927, shown for hotspot definitions: heat-source (dash-dot), 90% fusion (solid), 2keV (dotted) and density-average (dashed).

work for the same four definitions shown in figure 4.13. The heat source hotspot by definition has the strongest thermal conduction losses, and due to its small radius, does not encompass the entire radiation loss region, as seen in figure 4.13. As a result, the comparatively low radiation loss is comparable in magnitude to the thermal losses.

The alpha-heating contributions seem to increase with relative hotspot size, as expected from the inclusion of more of the radial profiles. The magnitude of the thermal conduction contributions decrease with relative hotspot size, as the larger hotspots incorporate more and more of the region of absorption of thermal heat flow. The magnitude of the radiative contributions increase at first, as the transition between radiative emission and absorption lies around $5 - 8\mu m$ beyond the transition for thermal conduction. This results in a maximal radiation loss for the 2keV hotspot, which has a relative thermal conduction loss of $\sim 1/2$. The density-average hotspot has a similar thermal conduction to radiation peak loss ratio of around half, while for the fusion hotspot this is $\sim 2/3$.

Figure 4.15 shows the burn-averaged ion temperatures and densities for the same hotspot definitions as above. We can see that the larger hotspot definitions result in higher burn-averaged ion temperatures and densities, with the 2keV and density-average hotspots consistently around $300 - 400eV$ hotter than the smallest heat-source hotspot. The burn-averaged density of the heat source hotspot is almost 30% smaller than for the density-average hotspot, which is the largest and therefore includes more of the shell. Note that while the intuitive mass-averaged ion temperatures would decrease with increasing hotspot definition radius, the fusion-weighting in the burn-averaged quantity results increase with increasing hotspot definition radius, as these larger hotspot definitions include more fusing material.

Figure 4.16 shows the hotspot radii for various hotspot definitions as a function of time for the robustly igniting $S = 1.0$ N130927 simulation, as figure 4.12 showed for
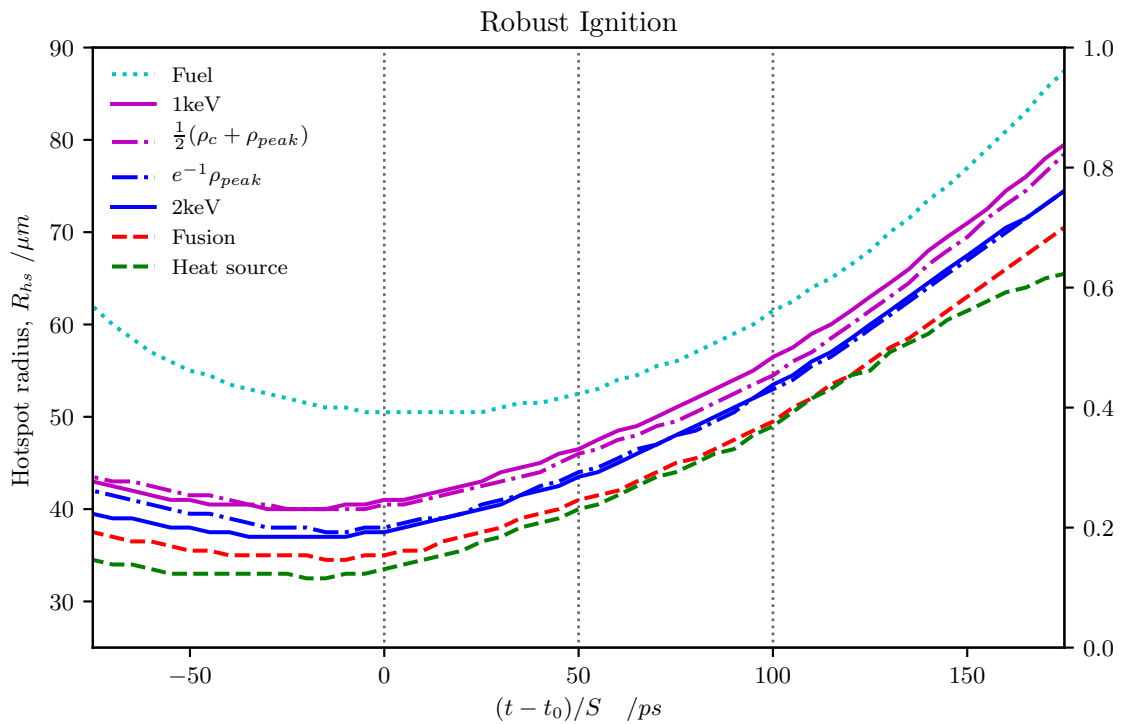
Figure 4.16: Hotspot radii against time for the scale $S = 1.0$ simulation of N130927 in the robust ignition regime in the same format as figure 4.12. The definitions included are the heat source (green dashed), 90% fusion volume (red dashed), 2keV (blue solid), 1keV (magenta solid), density e-fold ($e^{-1}\rho_{peak}$) (blue dash-dot), and density-average ($\frac{1}{2}(\rho_0 + \rho_{peak})$) (magenta dash-dot). The edge of the fuel is also shown (cyan dotted).
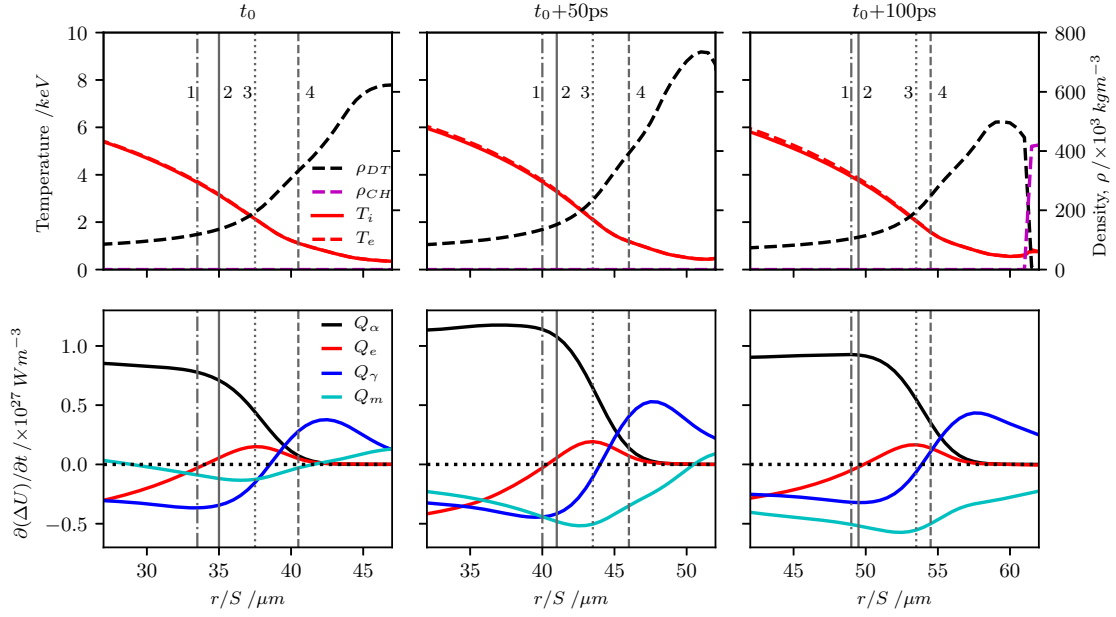
Figure 4.17:   An enlarged version of the radial profile time-evolution for the robust igni-
tion regime from figure 4.9, focussed on the hotspot-shell boundary region in
the same format as figures 4.13. Again the vertical lines indicate the hotspot
boundaries, for the following hotspots: (1) heat-source (dash-dot), (2) fusion
(solid), (3) $2keV$ (dotted) and (4) density-average (dashed).

the $S = 0.9$ self-heating case. Similarly, figure 4.17 shows the density, temperature and
power balance contribution radial profiles, zoomed in around the edge of the hotspot.

The various definitions remain similar to the self-heating case, except for the heat-
source hotspot, which is larger and more comparable in size to the fusion hotspot.
This is due to the higher alpha-heating levels in this regime producing higher central
temperatures. The 1keV and the density-average hotspots remain similar, as do the
2keV and the density e-fold hotspots. The range of the hotspot radii stays at around
$7 - 8\mu m$ as for the self-heating regime, but is small proportially due to the larger
hotspot sizes.

Figure 4.18 shows the volume-integrated hotspot power balance contributions for the
heat-source, 90% fusion, 2keV and density-average hotspots, as in figure 4.14 but for the
robustly igniting $S = 1.0$ regime. In all of the definitions shown, the radiative losses are
roughly comparable to the thermal conduction losses for this regime. The peak alpha-
heating contributions are $\sim 15\%$ lower for the smaller heat-source and fusion hotspots,
compared to the 2keV and density-average hotspots.

The comparison of burn-averaged ion temperature and density between hotspots for
the $S = 1.0$ robust ignition scenario shows that the variation in temperature decreases
to a gap of around $250 - 300eV$ in this stronger regime, compared to the self-heating
regime's $300 - 400eV$. Similarly, the difference in burn-averaged density between the
density-average and the heat source hotspots reduces from $20-30\times10^3 kgm^{-3}$ to around
$10 - 20 \times 10^3 kgm^{-3}$, or a $\sim 20\%$ difference.
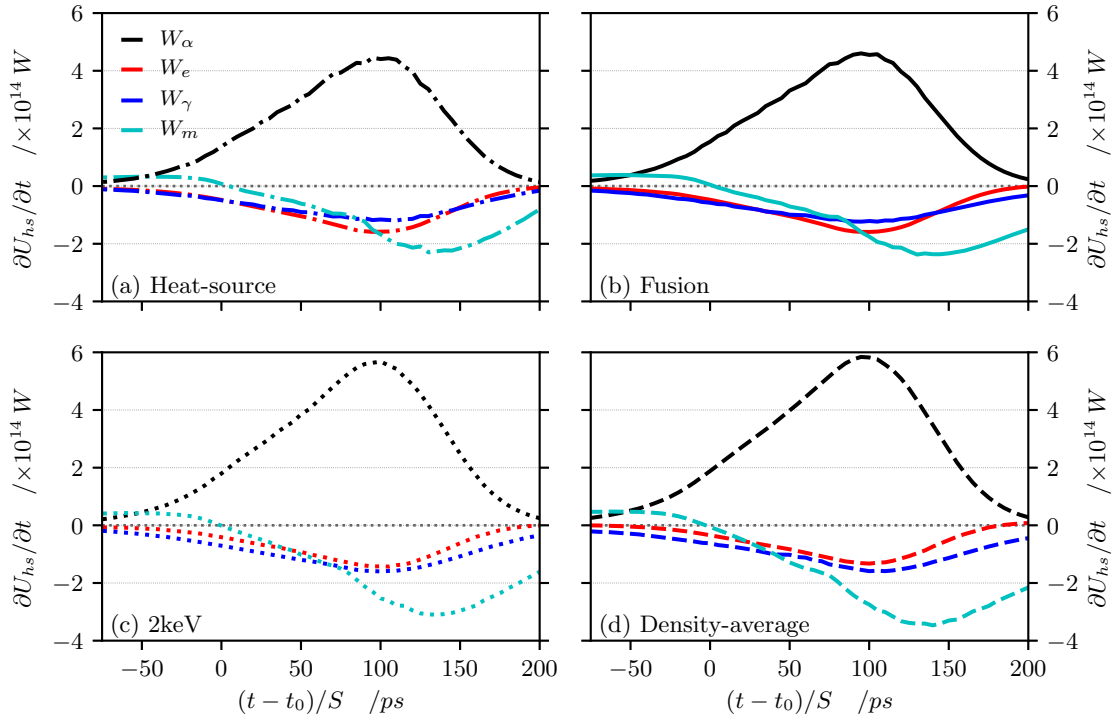
Figure 4.18: Volume-integrated contributions to the hotspot power balance for the robust ignition $S = 1.0$ scale simulation of N130927, in the same format as figure 4.14.
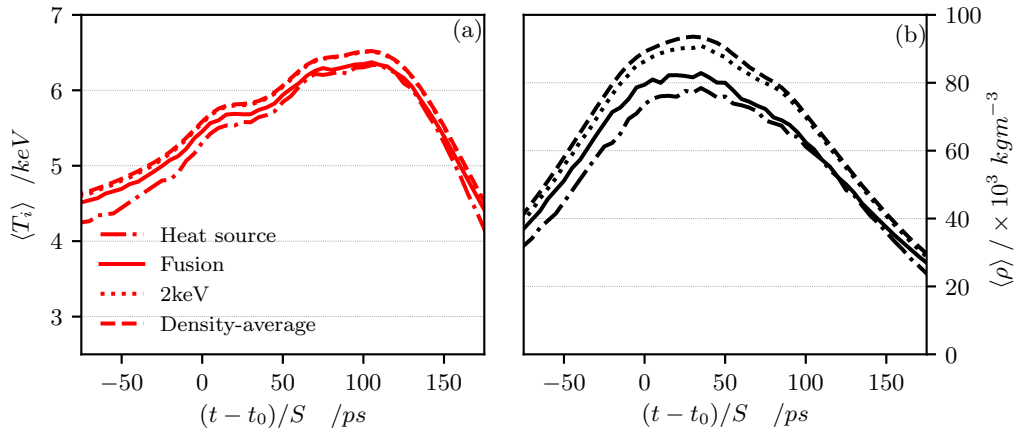
Figure 4.19: Burn-averaged (a) ion temperature, $\langle T_i \rangle$ and (b) density, $\langle \rho \rangle$ for the robust ignition $S = 1.0$ scale simulation of N130927, shown for hotspot definitions: heat-source (dash-dot), 90% fusion (solid), 2keV (dotted) and density-average (dashed).
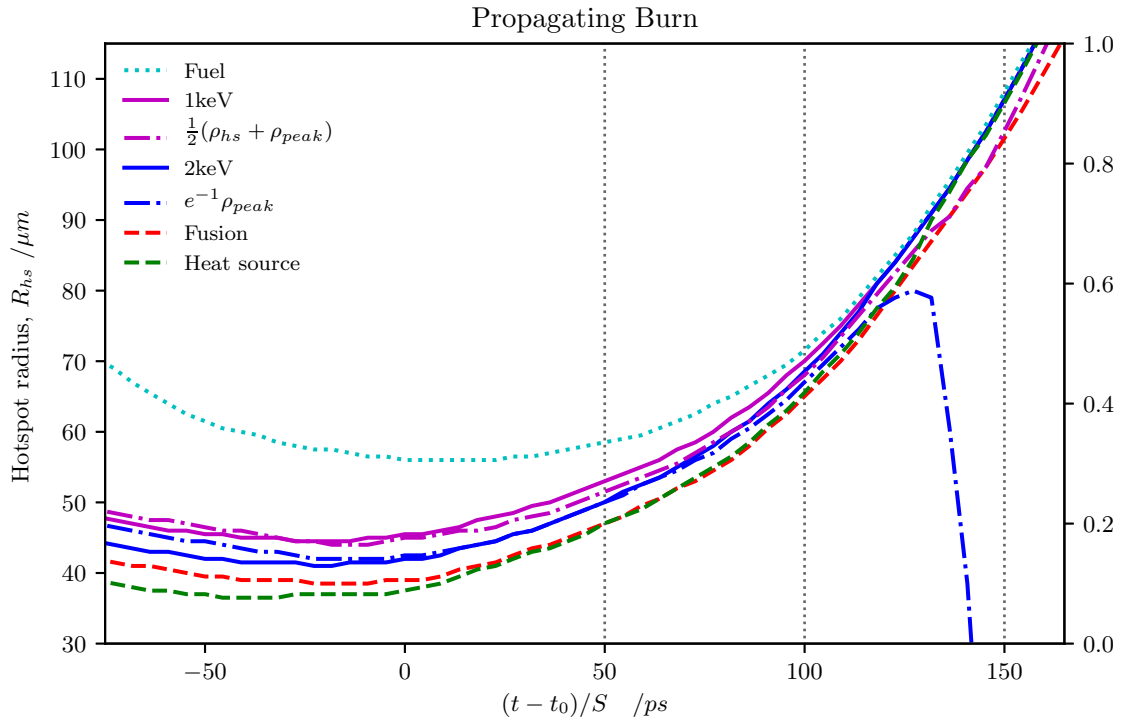
Figure 4.20:   Hotspot radii as a function of time for the scale $S = 1.1$ simulation of N130927 in the propagating burn regime, in the same format as figure 4.12. The definitions included are the heat source (green dashed), 90% fusion volume (red dashed), 2keV (blue solid), 1keV (magenta solid), density e-fold ($e^{-1}\rho_{peak}$) (blue dash-dot), and density-average ($\frac{1}{2}(\rho_0 + \rho_{peak})$) (magenta dash-dot), as well as the edge of the fuel (cyan dotted).
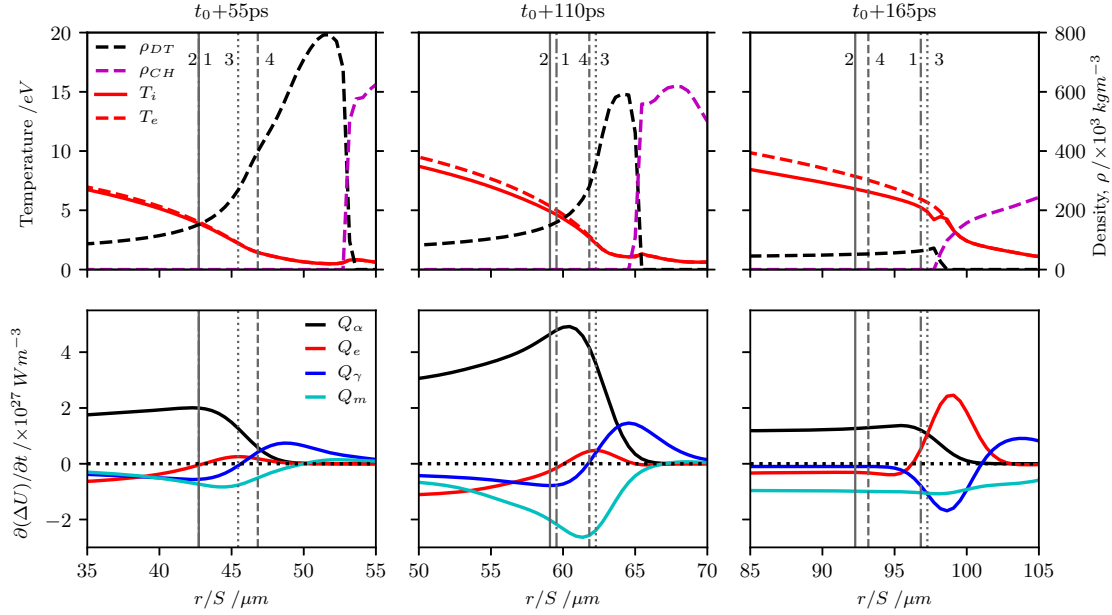
Figure 4.21: An enlarged version of the radial profile time-evolution for the propagating burn regime from figure 4.9, focussed on the hotspot-shell boundary region in the same format as figures 4.13 and 4.17. Hotspot boundaries are shown for the following hotspots: (1) heat-source (dash-dot), (2) fusion (solid), (3) $2keV$ (dotted) and (4) density-average (dashed).

Figure 4.20 show the hotspot radii time-variation for the $S = 1.1$ propagating burn scenario. As the hotspot ignites and begins to expand through burn propagation into the shell, the various definitions are compacted into a much smaller range. This is due to the increased heat flow, mass ablation and pressure reducing the temperature and density gradient scale-lengths (as mentioned in the previous section, section 4.3). The resultant size gap in hotspot radii varies from $\sim 10\mu m$ before stagnation (similar to self-heating) down to $\sim 5\mu m$. Notice that as the hotspot ablates mass away from the shell and reduces the peak density, the density e-fold ceases to be a useful definition, with the radius falling off towards 0 even as the hotspot expands and continues to heat up. Within this range, the positions of various hotspot definitions relative to one-another change a lot, with each definition expanding at a different rate. Figure 4.21, analogous to figures 4.13 and 4.17, shows the positions of different hotspot definition boundaries in relation to the radial profiles of the $S = 1.1$ propagating burn case, as shown in figure 4.9. The variation in hotspot expansion rate for different definitions is clear, with the heat source (1) and 2keV (3) hotspots increasing in radius faster than the fusion hotspot (2) or the density-average (4).

Figure 4.22 shows the volume-integrated hotspot power balance contributions for the heat-source, 90% fusion, 2keV and density-average hotspots for the propagating burn $S = 1.1$ regime, analogous to figures 4.14 and 4.18. The thermal conduction losses are greater than the radiative losses for all definitions here, while the mechanical work losses peak at around 50% of the peak alpha-heating contribution, similar to for the
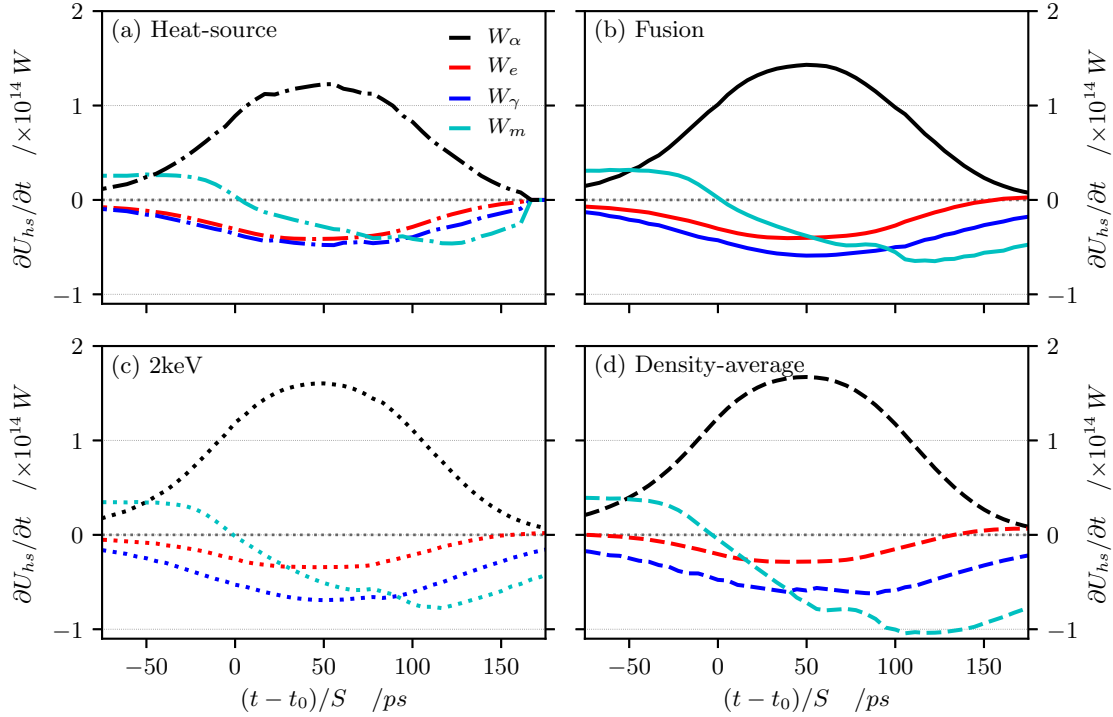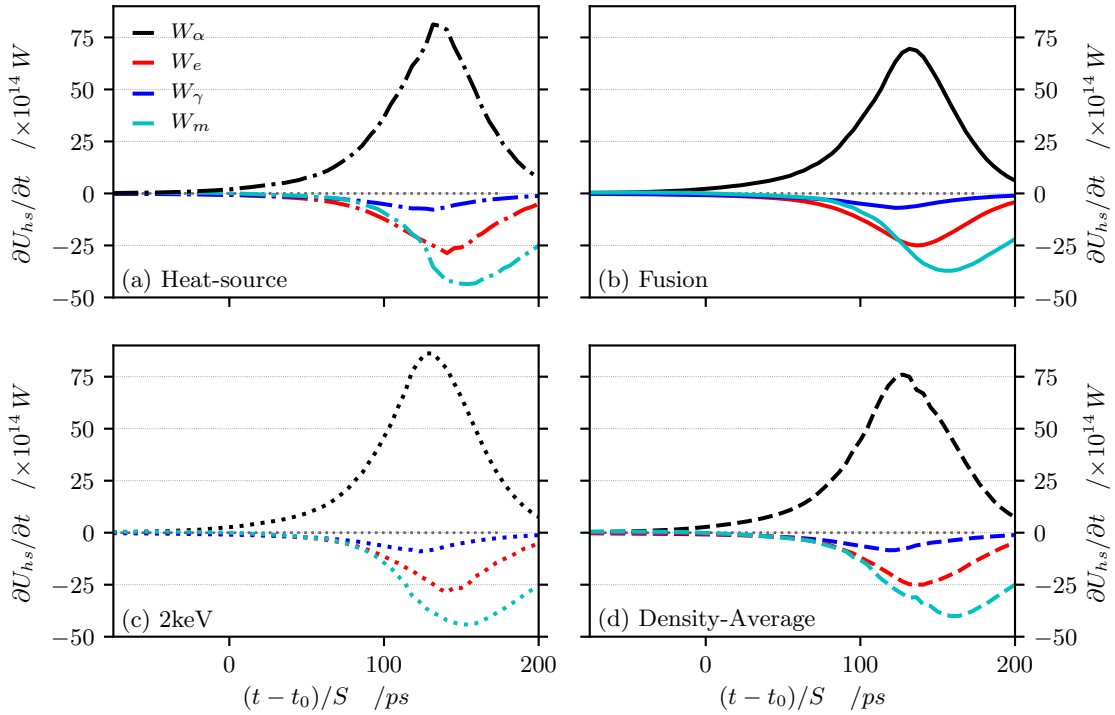
Figure 4.22: Volume-integrated contributions to the hotspot power balance for the propagating burn $S = 1.1$ scale simulation of N130927, in the same format as figures 4.14 and 4.18.

other regimes.

Considering figure 4.23, we can examine the variations in calculated burn-averaged temperature and density for this regime. The temperature variation reduces to less than $100eV$, which combined with the increased magnitude of temperature values equates to a $\sim 1\%$ variation between definitions. The density variation also decreases relatively to weaker alpha-heating regimes, to around $10\text{-}15 \times 10^3 kgm^{-3}$, around 20%. The 4 hotspot definitions seem to coalesce roughly into two positions, which is clearly shown in figure 4.23b. These density differences become negligible at $t_0 + 125ps$, as the expansion of the capsule becomes the dominant factor in determining the hotspot density.

4.5 Summary

In this chapter, we have explored the hotspot behaviour in a 1D spherical geometry. Work done by K. McGlinchey [130] to ensure accuracy of the base implosion characteristics (particularly in the ablation and implosion phases) through comparisons with simulations and experimental data from LLNL was outlined. This included the adjustment of radiation drive characteristics to match shock timings.

The radiative loss properties of the hotspot were explored by comparing two models, the full radiation transport model and an "elementary emission" model in which the radiative emission from the plasma is calculated based on its opacity; the elementary
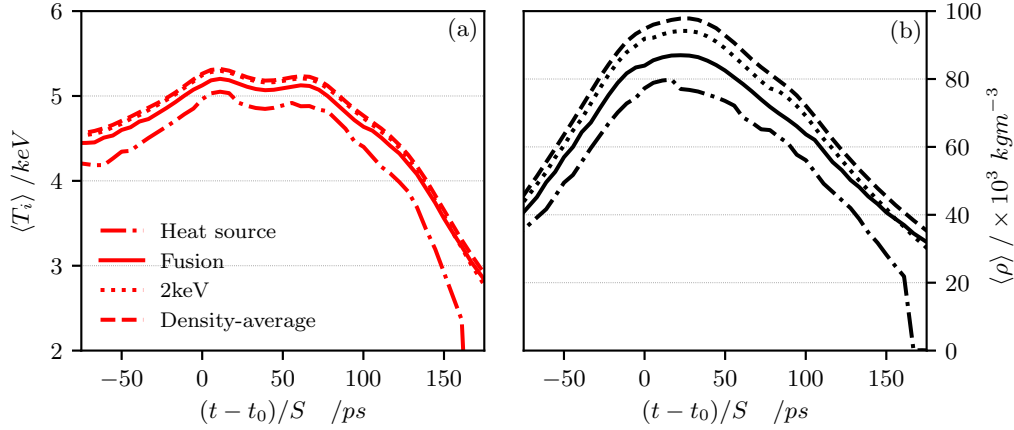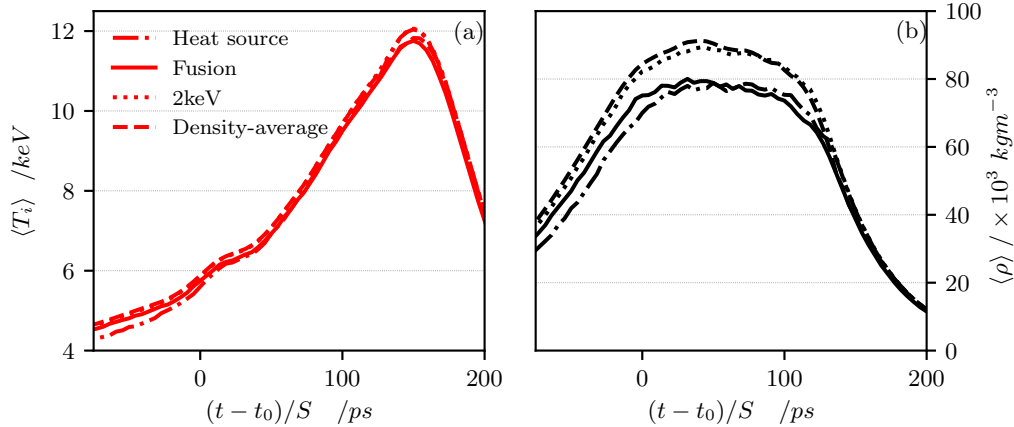
Figure 4.23: Burn-averaged (a) ion temperature, $\langle T_i \rangle$ and (b) density, $\langle \rho \rangle$ for the propagating burn $S = 1.1$ scale simulation of N130927, shown for hotspot definitions: heat-source (dash-dot), 90% fusion (solid), 2keV (dotted) and density-average (dashed).

emission model asymptotes to black-body emission and bremsstrahlung emission in the optically-thick and -thin regions respectively. Significant disparities were found between the models, with the elementary emission model hotspot emitting $\sim 3 \times$ more radiation than the full transport model hotspot. The hotspot self-absorption reduces the radiative emission from the hotspot by a factor of $\sim \frac{1}{2}$. In addition, the absorption of radiative energy in the shell is significant, and affects the mass ablation and scale lengths of the density- and temperature-gradients. We therefore expect the transport of radiation to play a non-negligible role in the evolution of the hotspot.

1D simulations based on the NIF Highfoot implosion N130927 were hydrodynamically scaled in order to explore different regimes of alpha-heating. The separate mechanisms of yield amplification between the robust ignition and propagating burn regimes were demonstrated. In the former case, yield amplification occurs due to an increase in the hotspot temperature only, while in the latter case it is also due to an increase in the total hotspot mass undergoing fusion reactions.

In the self-heating regime, the alpha-heating contribution to the power balance is unable to compensate for the power losses, such that $dW_{net}/dt < 0$ at peak compression. In the robust ignition regime, the alpha-heating contribution increases at a comparable rate to that for the combined losses and so $dW_{net}/dt \sim 0$. This results in the temperature of the hotspot being maintained (but not really increasing) even as the capsule expands, although there is little ablation of mass into the hotspot. In the propagating burn regime, the alpha-heating contribution significantly outweighs the combined losses, resulting in $dW_{net}/dt > 0$. The temperature of the hotspot continues to increase after peak compression, and the significant mass ablation from the hotspot heat flow into the shell means the hotspot density decreases at a much slower rate and the hotspot mass increases significantly.

It is worth emphasising the importance of the time-evolution of the hotspot within

these regimes. In order for the hotspot to propagate a deflagration wave into the shell, it must first undergo the characteristic temperature increases of robust ignition. As this occurs, it must also be adequately confined by the shell for heat flow to ablate significant material into the hotspot. The robust ignition regime effectively encompasses capsules that are igniting in the sense of producing enough fusion to have significant temperature increases, but are inadequately confined by the surrounding shell. Propagating burn requires a hotspot to be able to reach robustly igniting conditions, and also have sufficient confinement from the shell to develop a propagating burn wave.

We note that these regimes are also effectively characterised through Cheng *et al.*'s [87] and Springer *et al.*'s [85] prescriptions of hotspot ignition. Compared to Cheng's prescription [87, 125] that ignition effectively requires the heating timescale of the hotspot, τ_H, to be shorter than the confinement timescale of the hotspot, τ_C, the three alpha-heating regimes identified in this work can effectively be categorised into $\tau_H > \tau_C$, $\tau_H \sim \tau_C$ and $\tau_H < \tau_C$ for self-heating, robust ignition and propagating burn respectively. Similarly, considering Springer *et al.*'s [85] criterion that ignition requires both $dT/dt > 0$ and $d^2T/dt^2 > 0$, our regimes exhibit the characteristics of $d\langle T_i \rangle/dt < 0, dW_{net}/dt < 0$; $d\langle T_i \rangle/dt \sim 0, dW_{net}/dt \sim 0$; and $d\langle T_i \rangle/dt > 0, dW_{net}/dt > 0$ for self-heating, robust ignition and propagating burn respectively.

The alpha-heating deposition is roughly uniform across the hotspot in the self-heating regime, but transitions towards a Bragg-peak-like profile for stronger alpha-heating regimes. This is due in part to the higher temperature of the hotspot increasing the "transparency" of the hotspot to alpha-particles, and in part due to an increase in the number of alpha-particles being produced in close proximity to the shell. The electron thermal conduction deposits heat at a similar depth into the shell to the penetration depth of the alpha-heating deposition, with radiation being absorbed further out. Stronger levels of alpha-heating also drive stronger thermal conduction heat flows due to the boost in temperature, while the net radiation lost from the hotspot does not increase as quickly with strength of alpha-heating.

The alpha-heating has a significant impact on the hydrodynamics of the capsule implosion, owing to the non-linear feedback loop. The increases in hotspot temperature and pressure drive faster hotspot re-expansion, while the heat flow from the hotspot produces an ablative force and further compresses the shell. The temperature and density gradients between the hotspot and the shell are larger, resulting in significantly more distinct hotspot and shell regions than in a non-igniting scenario. This is in part because of the higher peak temperature and density in the hotspot and shell respectively, and in part due to the shell ablation and hotspot pressure.

We have explored ways in which the hotspot can be defined, and note that defining the hotspot by a particular density or temperature value, while convenient, has less physical reasoning behind it. The *ad hoc* nature of the value selection does not adapt well to different burn regimes — for example, while using the $1keV$ contour may make sense in the self-heating regime, the stronger heating in the robust ignition and propagating burn regimes pushes this contour much further out into the shell and ablator region. A

more conservative contour such as the $2keV$ contour is more robust to this, although still not dynamically adjusted to the strength of the alpha-heating in the hotspot. The experimentally measurable nature of the 17% neutron contour makes it uniquely useful, although the conservative self-similar neutron emission profile in the earlier calculation (in section 4.4.1) would give a hotspot that contained a much smaller proportion of the fusing material than many of the other hotspots considered here. However, stronger levels of alpha-heating would shift the neutron emission profile closer to the uniform emission, and therefore the 17% neutron contour hotspot would encompass more of the fusing material. The heat-source hotspot is intuitively very easy to understand, and by definition adjusts dynamically to account for stronger alpha-heating levels. The fusion volume is also dynamic in accounting for stronger alpha-heating levels, but fails when the entirety of the capsule is burning (and therefore ought to be considered as a part of the hotspot). Such definitions may be more involved to calculate and thus it is unclear that any particular definition is more correct than any other.

In general, we see that the variation in quantities such as the hotspot burn-averaged temperature and density, decreases with the increasing strength of the alpha-heating regime. In effect, different definitions of the hotspot include a different proportion of the boundary region between what is clearly definable as the hotspot and what is clearly definable as the shell. Thus, as a result of the hotspot becoming hotter and larger, the boundary region constitutes a smaller proportion of the entire hotspot, and therefore contributes less to the overall hotspot parameters. In addition, the boundary region decreases in width due to ablation and the increased pressure, and therefore there is less scope for variation in hotspot definitions. This can also be seen as the hotspot (as a region) becoming more distinct as the gradient scale-lengths decrease and the boundary between the hotspot and the shell becomes clearer. Conversely, it is at the weaker levels of alpha-heating that the uncertainty in hotspot quantities due to the hotspot definition are greatest, not only proportionally but in an absolute sense as well. Indeed, this is where any uncertainties stemming from hotspot definitions are the most important, as this is the regime which is currently the most experimentally relevant. It is clear that when we compare hotspot quantities, whether between experiments or simulations, it is necessary to know how the hotspot is defined, in order to accurately understand whether any differences are merely due to definition or in fact due to some physical effect.

5 3D Idealised Perturbations

In this chapter we investigate the behaviour of hotspots under the influence of idealised perturbations using 3D simulations. In the previous chapter, we examined the hotspot evolution towards igniting and burning in a 1D scenario by considering the time variation of the hotspot power balance, and the influence of alpha-heating. However, as noted in chapter 1, the capsule implosion is rarely, if ever, 1D, and therefore the consideration of higher-dimensional effects is required. The simplified context of idealised perturbations allows for a thorough investigation of their impact on the performance of a capsule, including the mechanisms of degradation and potential burn truncation. First, the initialisation procedure of 3D simulations is explained in section 5.1. We then explore two scenarios: in section 5.2, we examine the impact of a cold, dense spike from the shell impinging into the hotspot during its formation, and the subsequent impact on ignition. In section 5.3, we explore the impact of a region of reduced shell areal density (and therefore reduced confinement), and how it might affect the burn propagation.

5.1 Initialisation

In this section, the procedure for initialising our 3D simulations is explained. Due to the increased computational demand of 3D simulations compared to 1D simulations, a number of approximations have to be made in order to explore the 3D physics of the stagnation and burn phases of an ICF implosion.

5.1.1 1D-3D data remap

The first of these involves using radial 1D simulation data as a base for the 3D simulations. Running start-to-finish 3D simulations of ICF capsules with Chimera at a resolution which accurately resolves the physics in phases of the implosion would require an impractical amount of computational resources. Therefore, we instead remap 1D radial simulation data onto a 3D Cartesian grid at the time of peak radiation temperature, with cells being merged during said remap into a lower resolution grid. Even with a reduced simulation domain size, these 3D simulations need to be run at a reduced resolution owing to the N^3 scaling. The impact of regridding high-resolution data onto a lower resolution in a stagnation-phase reinitialisation was also investigated by K. McGlinchey [130]. Regridding to a lower resolution of $2\mu m$ produced a yield reduction of 5% compared to the nominal $0.25\mu m$ case, while a full simulation at the lower resolution reduced the yield by 20%. This indicates that using the high-resolution 1D simulation to more accurately capture the ablation phase physics and length-scales,

and then using lower resolutions from the stagnation phase onwards produces reasonable agreement with the using high-resolution the entire way through. We can increase the velocity by a small percentage at the point of reinitialisation in order to compensate for decrease in yield from remapping to lower resolutions.

5.1.2 Radiation group structure

We also introduce the concept of hydrodynamic decoupling [130], in which the radiation absorbed by the ablator at times beyond the peak radiation temperature is unable to affect the hydrodynamics of the hotspot formation during the stagnation phase, due to the sufficiently high implosion velocity of the shell. Following this, the 3D simulations are restarted at the time of peak radiation temperature without an external radiation source, which also allows us to use a reduced number of radiation groups. The level of atomic detail given by the structure of the 54-group radiation transport is necessary to ensure accuracy of radiative deposition during the early phases of the implosions. However, without an external radiation source, the detail of the atomic data required is reduced to only what is necessary to model the radiation transport within the hotspot. This fully ionized, pure DT plasma is much simpler to model than the complex radiative dynamics of partially ionized carbon involved in the ablation phase. Hence, the reduced 10-group structure allows for accurate modelling of the hotspot radiative losses while also reducing the computational demand of the radiation transport by a factor of 5. Both the hydrodynamic decoupling approximation and the reduced radiation group structure were compared to the full scenario by K. McGlinchey [130], showing good agreement in the time and magnitude of the peak implosion velocity, as well as in the density and temperature profiles at the time of peak neutron output.

5.1.3 Perturbations

Perturbations are applied in the velocity field after the remap, using Layzer's approximate analytic treatment [186] of the Rayleigh-Taylor instability, following S. Taylor's methodology [91]. We briefly summarise the methodology here, a full description of which can be found in S. Taylor's thesis work [91].

Layzer's cylindrical model for the velocity field is given as:

$$u_z = F(t)\mathrm{e}^{-|z|}J_0(\rho) \tag{5.1}$$

$$u_\rho = F(t)\mathrm{e}^{-|z|}J_1(\rho) \tag{5.2}$$

where $\mathbf{u} = \mathbf{u}(z, \rho)$ is the velocity field of the Rayleigh-Taylor bubble, z is the cylindrically-symmetric axis of the spike, ρ is the radial distance from the z-axis, $F(t)$ is some time-dependent amplitude, and J_0, J_1 are the zeroth- and first-order Bessel functions respectively. When applying these to our scenario, the cylindrical axis z becomes the capsule radial axis r. The perturbing velocity field is applied on top of the existing velocity field at the point of maximum density gradient on the outer surface of the cap-

sule in order to emulate the feed-through mechanism of ablation-phase perturbations. S. Taylor showed that applying the perturbation in this manner reproduces the expected behaviour of the Rayleigh-Taylor instability in the deceleration phase [91].

Although in this chapter we apply the perturbing field at one point only, these perturbations can be applied in a quasi-spherically symmetric manner by using the vertices of a geodesic sphere generated from an icosahedron (a regular 20-faced polyhedron) [91], as used in chapter 6. This method generates an approximation to a sphere by subdividing the triangular faces of the icosahedron into more regular triangles, and projecting these vertices onto the surface of a sphere. The order of the geodisation determines the number of vertices used to approximate the sphere, and therefore the distance between each vertex. The wavelength of a given "mode" of perturbation is therefore determined by the order of the geodisation used to generate it, with higher orders producing smaller wavelengths.

Superposing multiple modes can generate multi-mode perturbations rather than single-mode. By rotating the co-ordinates of the vertices produced by different orders of geodisation through random angles, coherent interference between modes can be avoided. The perturbations can also be pointed slightly off-centre from the central point of the simulation to further avoid coherent interference, and to add non-spherical symmetry to the perturbation. For the multi-mode case, the amplitudes are randomised using a uniform distribution between 0 and the given amplitude, although not for the single-mode case to avoid unintended long-wavelength perturbations.

5.2 Single Spike

Perturbations which grow on the outer ablator surface can feed through to the dense shell and produce spike-like structures which intrude into the hotspot. Therefore, in this section, we investigate the impact of spike-like perturbations on the hotspot through the simplest case of a single, idealised spike. We explore the interaction of the perturbation with the hotspot heat flow, considering how it affects hotspot properties as well as the converse of how the heat flow affects the perturbation growth. This is investigated for the robust ignition regime, as we look to understand how the perturbation may prevent or otherwise impede ignition.

5.2.1 Setup

We use the 1D simulation of N130927 mentioned in section 4.1 as the base for this simulation, at scale $S = 1.0$. Using S. Taylor's Layzer-like perturbation methodology (section 5.1.3), a single-mode perturbation of $\lambda \sim 40 \mu m$ is applied at a radius of $304 \mu m$, with the capsule at a convergence ratio of ~ 3 at the point of reinitialisation. The velocity perturbation amplitude $\delta v / \langle v_{imp} \rangle \sim \frac{2}{3}$, where the shell-averaged implosion velocity $\langle v_{imp} \rangle \sim 370 - 380 km/s$. This generates a single, inward propagating spike of cold, dense shell. A small velocity multiplier of 2% is also applied to the velocity field
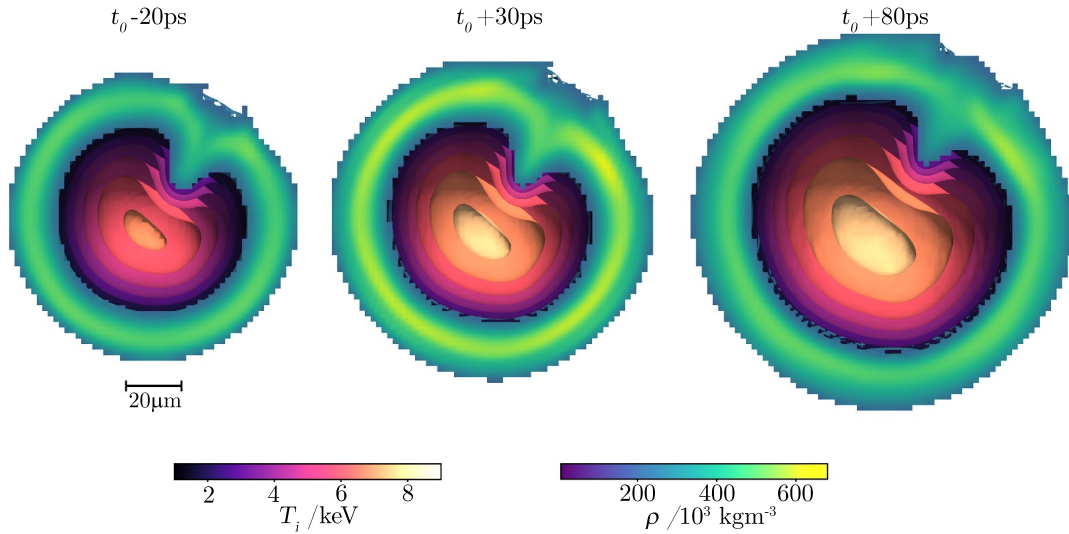
Figure 5.1: Hydrodynamic time evolution around peak compression of the single-spike implosion based on a scale $S = 1.0$ Highfoot N130927 shot, showing the ion-temperature contours of the hotspot, and a slice through the shell density. A velocity multiplier of 2% has been applied in order to account for the lack of radiation drive, and the effect of reduced resolution on the capsule performance.

throughout, in order to account for the reduction in yield stemming from the reduced resolution.

5.2.2 Results

Figure 5.1 shows the hydrodynamic time evolution around peak compression (at $t = t_0$) of the single-spike simulation, showing the 3D ion temperature contours in the hotspot, and a 2D central slice through the shell. The spike perturbation can be clearly seen penetrating the hotspot from the top right of the figure. Note that although the perturbation does not appear to intrude too deeply in density, the distortions in the ion temperature extend much deeper. This is because the variation in density within the hotspot is small relative to the variation between the hotspot and the shell, the majority of which occurs at the edge of the hotspot. By contrast, the majority of the temperature variation between the hotspot and the shell occurs within the hotspot itself.

Figure 5.2 shows the time evolution of radial profiles along the perturbed line-of-sight (LoS) directly into (solid) and the unperturbed LoS away (dashed) from the perturbation (along the axis of the spike), showing density and ion temperatures (top) and the hotspot power balance contributions (bottom), including alpha-heating, thermal conduction, radiation and mechanical work. The time-steps shown are the same as in figure 5.1. The density and temperature profiles are clearly compressed towards the centre for the perturbed LoS compared to the unperturbed LoS. The intruding spike reduces the hotspot radius along the line of sight by up to $20\mu m$.

The temperature drops rapidly from the centre for the perturbed LoS, in a slightly
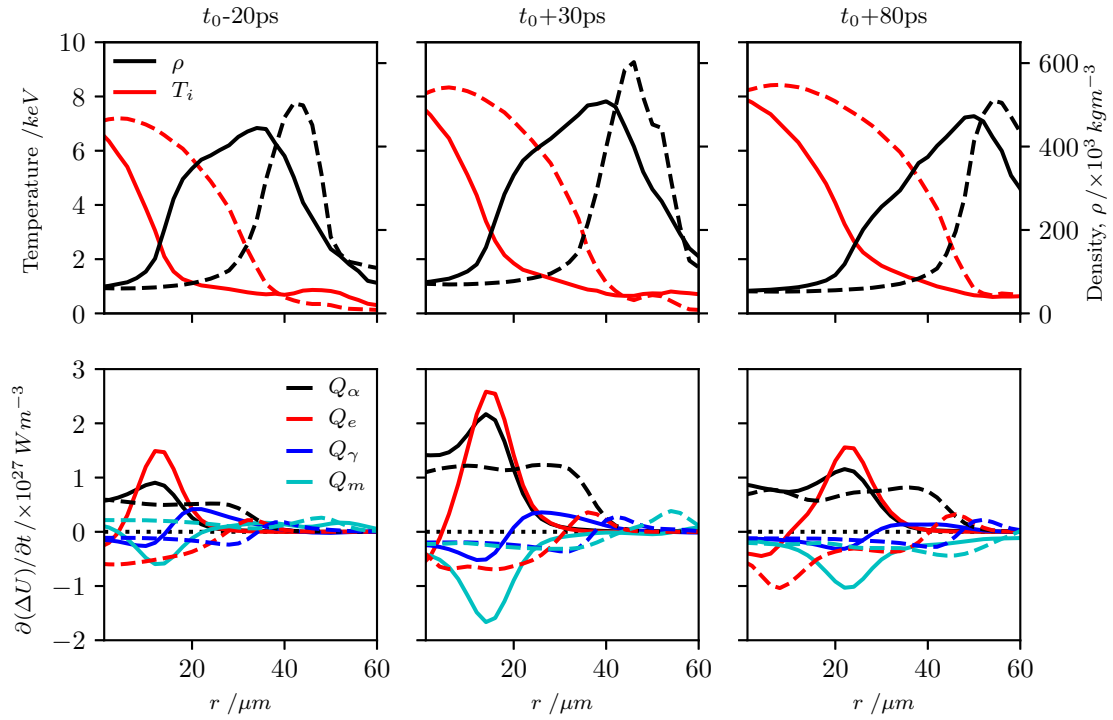
Figure 5.2: The time evolution of radial profiles directly into (solid) and away from
(dashed) the perturbation spike, showing (top) density (black) and ion temperature (red), and (bottom) the hotspot power balance contributions from
alpha-heating, Q_α (black), electron thermal conduction, Q_e (red), radiation,
Q_γ (blue) and mechanical work, Q_m (cyan).
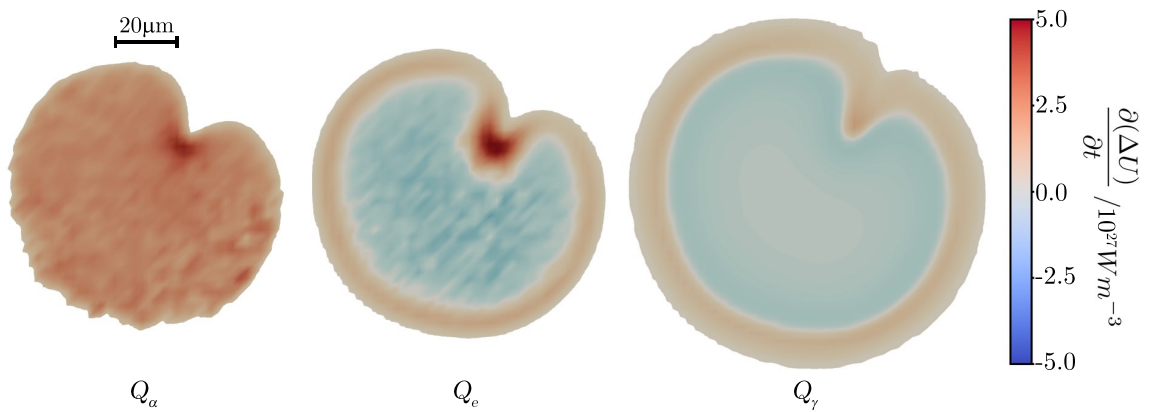
Figure 5.3: 2D central slices showing the power deposition maps for the single spike at
$t = t_0 + 30$ps for alpha-heating, Q_α (left), thermal conduction, Q_e (middle)
and radiation, Q_γ (right). All maps are shown on the same colour scale, with
blue indicating regions of power emission or loss, and red indicating regions
of power absorption or gain.
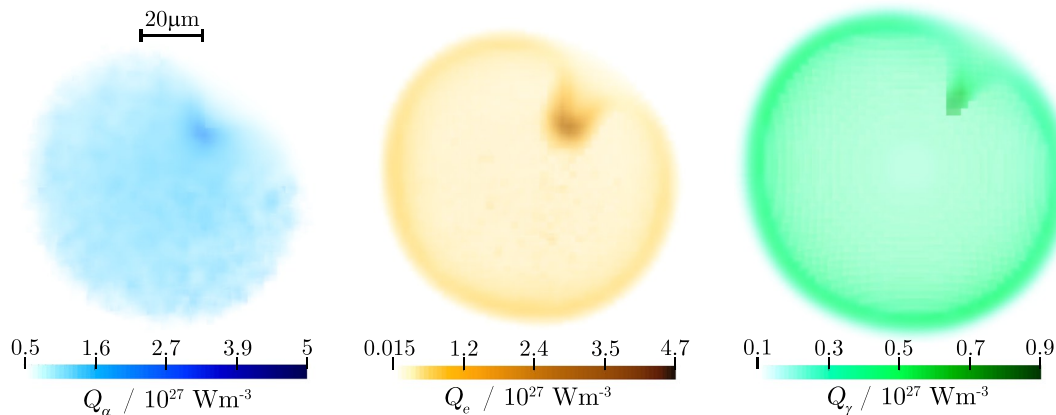
Figure 5.4: 3D volume plots showing the power absorption only (not emission) for the single spike at $t = t_0 + 30$ps for alpha-heating, Q_α (left), thermal conduction, Q_e (middle) and radiation, Q_γ (right). The darker regions indicate stronger absorption.

sharper manner than for the unperturbed LoS. The temperature gradient directly into the spike is slightly higher than for the unperturbed case, and the peak gradient also occurs closer to the centre of the hotspot and therefore at a higher temperature. The combination of this, along with the increase in surface area results in significantly more thermal conduction being absorbed into the spike region; the peak thermal conduction absorption is $\sim 5\times$ larger in the perturbed LoS than for the unperturbed LoS.

The alpha-heating into the spike peaks at $\sim 2\times$ that of the unperturbed LoS. Alpha-particles not encountering the perturbation still deposit the majority of their energy within the hotspot as they transit across it, leaving little to deposit into the shell — as in the case for the 1D scenario of robust ignition. However, the proximity of the cold dense spike to the centre of the hotspot results in a much larger flux of alpha-particles into the spike, which then deposit the majority of their energy in a very small region owing to the higher stopping power of the cold dense material. Note that these particles retain a significant proportion of their initial energy at the point of encounter with the spike.

The PdV expansion work also increases (more negative on the figure) around the spike, due to the stronger heat deposition (from thermal conduction and alpha-heating as above) producing stronger ablation and mass flow off the shell material into the hotspot. The radiative contribution into the spike is relatively similar to the unperturbed LoS, with the largest difference simply being the difference in radial extent due to the spike's penetration depth into the hotspot. At later times, the radiative absorption and the edge of the thermal conduction and alpha-particle heat flow have heated the material and expansion has lowered the density, such that the radiation lost from the hotspot and the radiation absorbed at the edge both decrease in magnitude.

These 3D effects are best illustrated by figures 5.3 and 5.4 rather than the 1D radial profiles of figure 5.2. Figures 5.3 and 5.4 show the alpha-heating, thermal conduction and radiative contributions to the power balance at $t = t_0 + 30$ps. Figure 5.3 shows

these contributions as a 2D slice representation of the deposition map in the plane of the perturbation, with all contributions shown on the same colour and spatial scale for ease of comparison. Figure 5.4 shows 3D volume representations of only the absorption component of the power balances, with the darker regions indicating stronger absorption. Each component is shown on its own individual scale.

From figure 5.3, the alpha-heating can be seen to be relatively uniform throughout most of the hotspot, agreeing well with the deposition maps seen in chapter 4. The solid angle around the spike has reduced alpha-heating deposition, which absorbs the majority of particles propagating in that direction and produces a shadow behind it. The extent of the alpha-heating is similar to that of the thermal conduction, which mostly deposits in a thin shell around the hotspot, as expected from the radial profiles of figure 5.2 and chapter 4. There is strong thermal conduction heat deposition into a broad region around the spike — this is due to the increase in surface area around the spike's distortion, in combination with an increased temperature gradient, and the position of the peak temperature gradient moving deeper into the hotspot (i.e. closer to the centre, and therefore to higher temperatures and thermal conductivities ($\kappa \propto T$)). Similarly, the radiative deposition is stronger only in a narrow region deep into the spike, but with only a small increase in magnitude. These depositions are cylindrically symmetric around the axis of the spike, as can be seen from figure 5.4. The shell-like nature of the thermal conduction and radiation heat flow depositions is clearly shown, with the regions of strong absorption around the spike from figure 5.3 also visible.

Overall, it seems that the thermal conduction is the most significant component in the power balance into the spike, while the alpha-heating is more important in the unperturbed shell regions. The increase in surface area has a larger effect on concentrating the thermal conduction heat flow out of the shell than on the alpha-heating heat flow. This is because the thermal conduction heat flows perpendicular to the temperature contours (and effectively, the surface of the hotspot), while the alpha-particle flux is net radially outwards, since the centre of the hotspot is the region of strongest production. The regions around the sides of the spike produce fewer alpha-particles than the centre, and therefore the sides of the spike do not see the same flux as the tip of the spike.

Figure 5.5 compares various hotspot volume-integrated properties over time for the single-spike simulation and a companion unperturbed, symmetric P_0 simulation, showing the burn-averaged density, $\langle \rho \rangle$, ion temperature, $\langle T_i \rangle$, mass as a percentage of the total fuel, M_V/M_{DT}, volume, V and surface area, SA, in addition to the overall burn history. Figure 5.6 shows the volume-integrated hotspot power balance contributions for the single-spike and companion symmetric simulations, with parts (b) and (c) showing zoomed in regions on the early-time power balance.

The perturbation reduces the yield by $\sim 40\%$, from 1.4×10^{17} to 8.1×10^{16}, as well as shifting bang time 27ps earlier. Within the hotspot, $\langle T_i \rangle$ diverges at roughly $t = t_0 - 50$ps, while the differences in the burn history are only visible later on. After peak compression, $\langle T_i \rangle$ flattens off quickly for the single-spike case, while it continues to increase for the unperturbed case. This is inextricably linked with the weaker alpha-
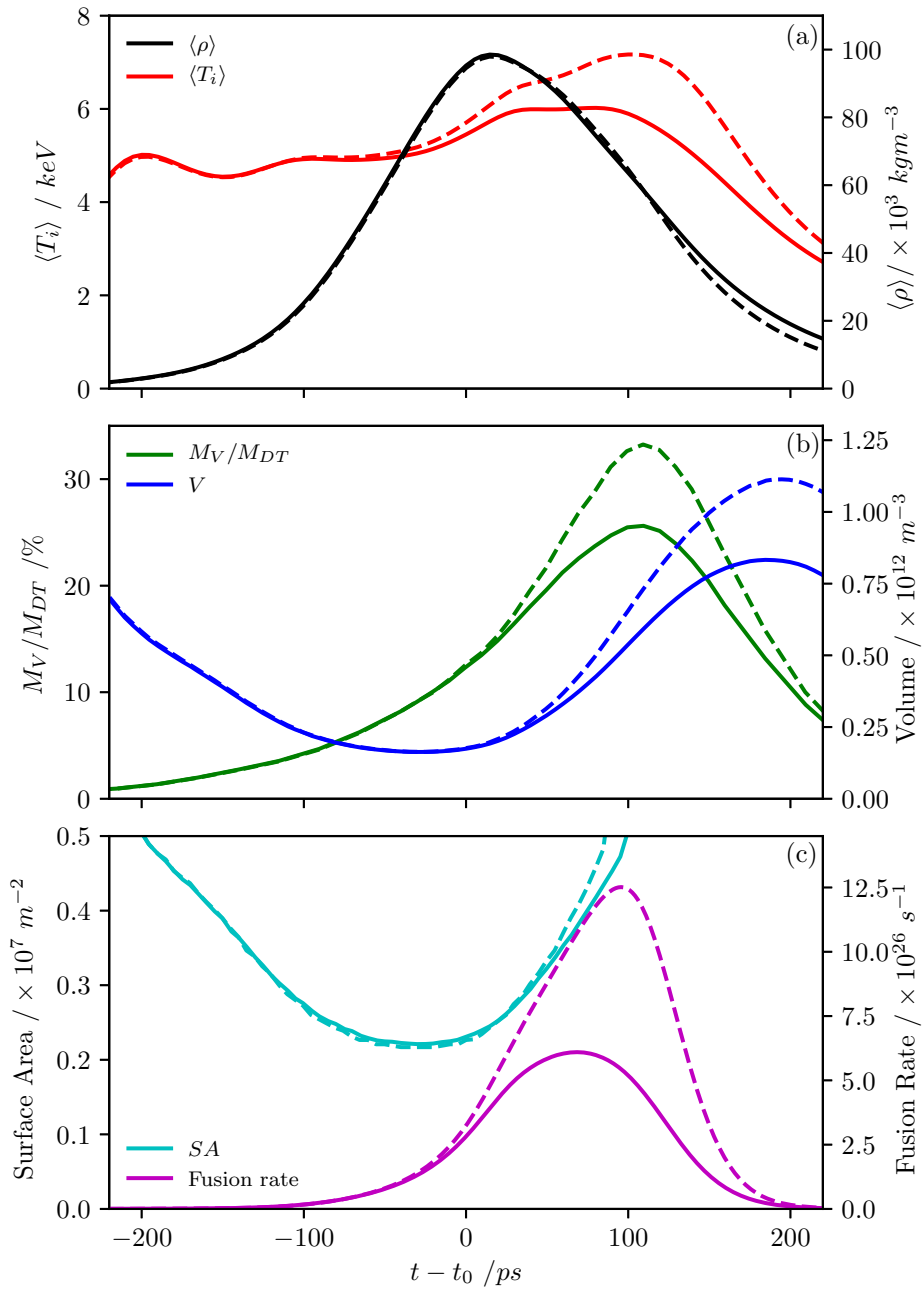
Figure 5.5: A comparison between the single-spike (solid) and the symmetric companion (dashed) simulations, showing: (a) hotspot burn-averaged density (black) and ion temperature (red); (b) hotspot volume (blue) and mass as a fraction of the entire fuel (green); and (c) hotspot surface area (cyan) and the neutron burn history (magenta).
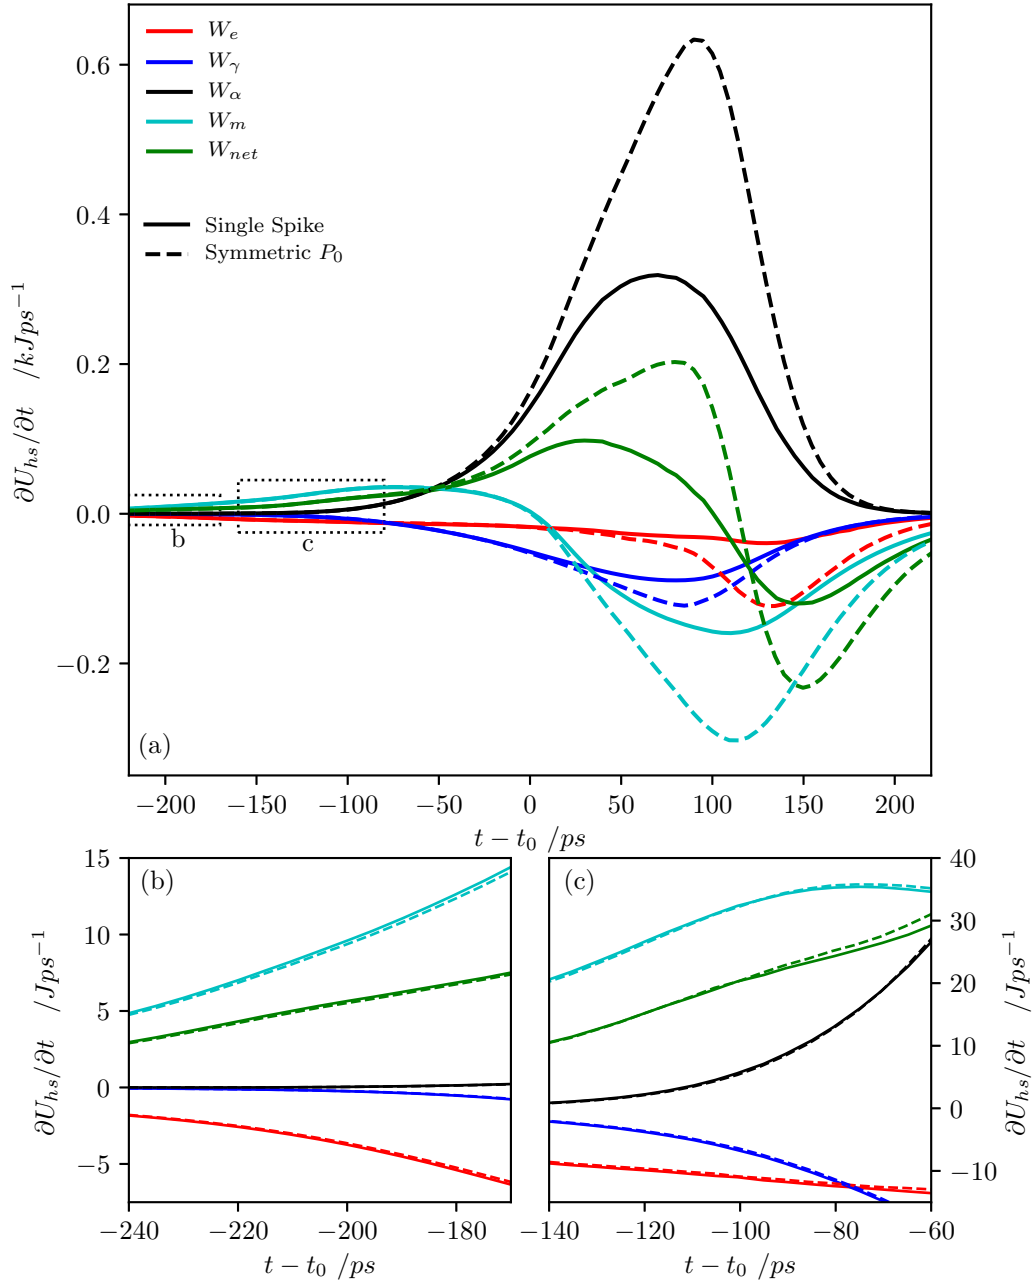
Figure 5.6: (a) Volume-integrated power balance contributions for the 2keV hotspot for the single-spike (solid) and companion symmetric (dashed) simulations, shown for alpha-heating, W_α (black), electron thermal conduction, W_e (red), radiation, W_γ (blue), mechanical work, W_m (cyan) and total, W_{net} (green). (b) and (c) are zoomed in on the regions indicated by dashed boxes in (a), showing the early-time variations during the deceleration phase.

heating.

The burn-averaged densities are almost identical between the perturbed and unperturbed simulations, with significant deviations only after the majority of burn has occurred. The single-spike density initially drops off slightly faster than for the symmetric case; somewhat counter-intuitively, the symmetric case ablates more material as the stronger ignition produces a stronger heat flow. Although material is ablated off the cold dense spike in the single-spike case, weakened alpha-heating bootstrap results in less shell ablation overall than for the symmetric case. The single-spike hotspot has both a lower fractional mass and a lower volume than the symmetric case, which matches the similar densities between the two scenarios. Around peak compression, the single-spike case has a higher surface area, but falls below that of the symmetric case as the symmetric case heats up more and expands faster after peak compression.

The weaker bootstrap in the single-spike scenario produces not only a lower net heating, but also weaker contributions from thermal conduction, radiation and PdV due to the lower temperature, volume and pressure in the hotspot respectively. The thermal conduction is particularly low for the single-spike case; this is due to the sensitivity of the hotspot boundary as defined by the 2keV contour, which includes a significant proportion of the spike and therefore reduces the calculated magnitude of the thermal conduction loss.

The early-time (before peak compression) differences in the power balance manifest themselves through lower hotspot temperatures as the shell stagnates. This is the most important time, as the conversion of mechanical work done during the stagnation to hotspot thermal energy is what determines the strength of the alpha-heating bootstrapping — effectively the 'spark'. The perturbation moves in faster than the rest of the shell and increases the PdV work early on ($t \sim t_0 - 200$ps). However, overall the symmetric case has a more synchronous shell stagnation, resulting in a higher peak PdV work which produces hotter temperatures than for the perturbed case (from $t \sim t_0 - 100$ps). These hotter temperatures result in stronger alpha-heating and therefore a stronger bootstrap.

We note that the yield degradation due to enhanced power losses in the hotspot after formation and yield degradation due to a weaker alpha-heating feedback loop from early-time variations in the conditions of the hotspot formation are distinct processes, but difficult to distinguish from one another. We can attempt to do so by considering the hotspot as two separate but connected parts, and comparing the properties of each of the perturbed and unperturbed 'hemispheres' (noting that the perturbed 'hemisphere' will not actually be hemispherical in shape). These comparisons are shown in figures 5.7 and 5.8. Figure 5.7 compares the hotspot burn-averaged density, burn-averaged ion temperature, mass fraction, volume, surface area and fusion rates, analogous to figure 5.5 but for the two hemispheres of the single-spike simulation rather than across two simulations. Similarly, figure 5.8 is analogous to figure 5.6, showing the volume-integrated power balance contributions for the two hemispheres.

The perturbed hemisphere has a lower temperature and higher density than the un-

Figure 5.7: A comparison between the hemisphere containing the perturbation (solid) and the hemisphere without the spike perturbation (dashed), showing: (a) hotspot burn-averaged density (black) and ion temperature (red); (b) hotspot volume (blue) and mass as a fraction of the entire fuel (green); and (c) hotspot surface area (cyan) and the DT fusion reaction rate (magenta).
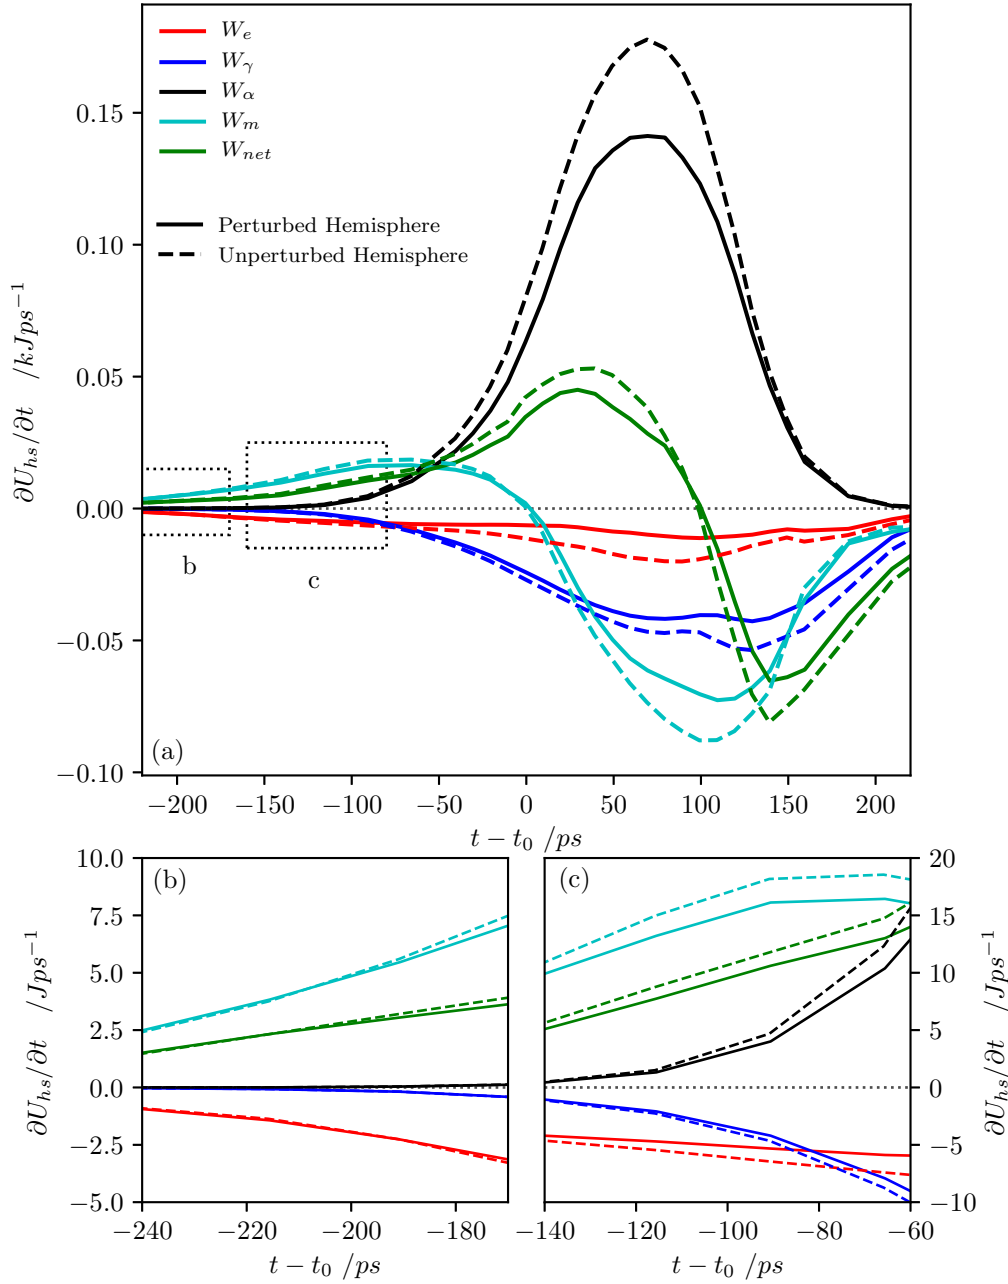
Figure 5.8: (a) Volume-integrated power balance contributions for the 2keV hotspot for the perturbed hemisphere (solid) and unperturbed hemisphere (dashed) simulations, shown for alpha-heating, W_α (black), electron thermal conduction, W_e (red), radiation, W_γ (blue), mechanical work, W_m (cyan) and total, W_{net} (green). (b) and (c) are zoomed in on the regions indicated by dashed boxes in (a), showing the early-time variations during the deceleration phase.

perturbed hemisphere. The higher density is likely due to the colder hemisphere being more compressible, which combined with the lower mass fraction also result in a lower volume and surface area. The reduced fusion rate also results from the lower temperature. The perturbed hemisphere has lower magnitudes for all of the contributions to the power balance, and also a lower net heating. Before $t = t_0 - 210\text{ps}$, more PdV work is being done on the perturbed hemisphere, but the unperturbed hemisphere passes it after this point. This is due to the poor synchronisation of PdV work from the shell due to the faster-moving spike, as mentioned earlier. The lower temperatures throughout the perturbed hemisphere also produce less alpha-heating in the region, and partially explain the lower thermal conduction and radiative losses, which are also affected by the lower volume and surface area. The spike's presence reduces the temperature of the entire hemisphere, through both the poor PdV work synchronisation and the weaker alpha-heating within the volume.

Contrary to any notion that a spike might feed extra fuel to the hotspot and allow it to burn more material, the cold, dense shell material transported towards the hotspot in this regime acts as a sink, drawing energy out of the hotspot and reducing the bootstrap heating. The spike enters the hotspot before the burn has had a chance to develop, and therefore both reduces the strength of the ignition 'spark' and impedes the development of a burn wave. It is perhaps possible that the spike might be able to feed the hotspot burn, but only in the difficult scenario of having ignited the hotspot before the spike can influence it. This would require the hotspot to begin propagating a burn wave before having reached peak compression.

5.3 Single Bubble

The confinement of a hotspot is determined by the areal density of the confining shell. Significant deviations from uniformity in the shell can therefore result in asymmetric confinement of the hotspot. Of particular interest is the possibility of confinement loss through a region of locally weak confinement in the shell, i.e. a region of particularly low areal density. In this section, we investigate the impact of a perturbation designed to produce a region of low areal density in the shell, and therefore result in the hotspot 'bubble' expanding faster into this region.

5.3.1 Setup

As in section 5.2, we again use the 1D data on N130927 as the base, again at scale $S = 1.0$. Here, instead of using the velocity field perturbation to generate a bubble, the Layzer perturbation structure is used to apply a reduction in the density to the shell. Using a velocity field to generate the bubble would effectively require a significant reduction in shell convergence due to a low implosion velocity, whereas the proposed method should result in the natural evolution of a bubble as the shell re-expands under the central hotspot pressure. The region with reduced density will have less inertia
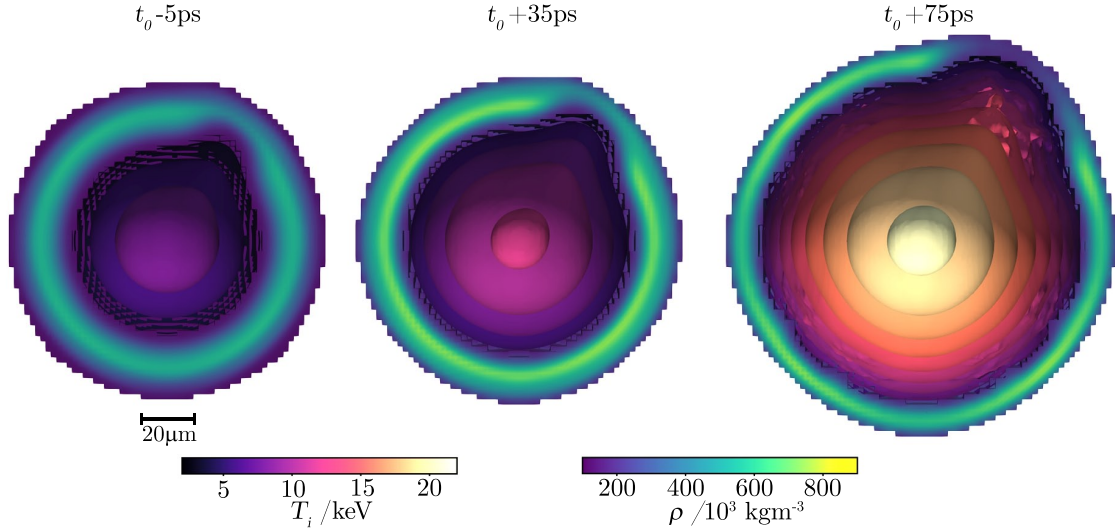
Figure 5.9: Hydrodynamic time evolution around peak compression of the implosion based on a scale $S = 1.0$ Highfoot N130927 shot perturbed with a density reduction to induce a bubble-like perturbation, showing the ion-temperature contours of the hotspot, and a slice through the shell density. A velocity multiplier of 5% has been applied in order to account for the lack of radiation drive, and the effect of reduced resolution on the capsule performance.

to confine the hotspot, and will expand faster, producing a bubble, or an 'aneurysm' [39, 187]. This method does not conserve the mass in the shell but simply removes it, whereas a more physical method would redistribute the mass from the region to other regions in the shell. We justify this method as simply a 'toy' model which is used to investigate the underlying physics, rather than as an accurate depiction of a particular experiment.

We wish to explore this perturbation for a hotspot with sufficiently high pressure to organically generate a bubble from the region of low areal density. A velocity multiplier of 5% is applied to the capsule in order to boost the implosion velocity, and therefore boost the capsule performance. In this simulation, we also explore the possibility of attempting to reach the propagating burn regime through simply increasing the implosion velocity of a design rather than undergoing full hydrodynamic scaling of entire capsule, as this effectively necessitates a larger driver.

5.3.2 Results

Figure 5.9 shows the hydrodynamic time evolution around peak-compression of the simulation with the single-bubble perturbation applied. As in figure 5.1, the ion temperature contours of the hotspot can be seen in the centre of the capsule, alongside a 2D slice through the centre of the density of the confining shell. The bubble can be seen forming already in the first time-step just before stagnation, with the velocity flow beginning to converge towards this region. This is because the reduced shell density of the region has less incoming ram pressure, and therefore stagnates earlier than the
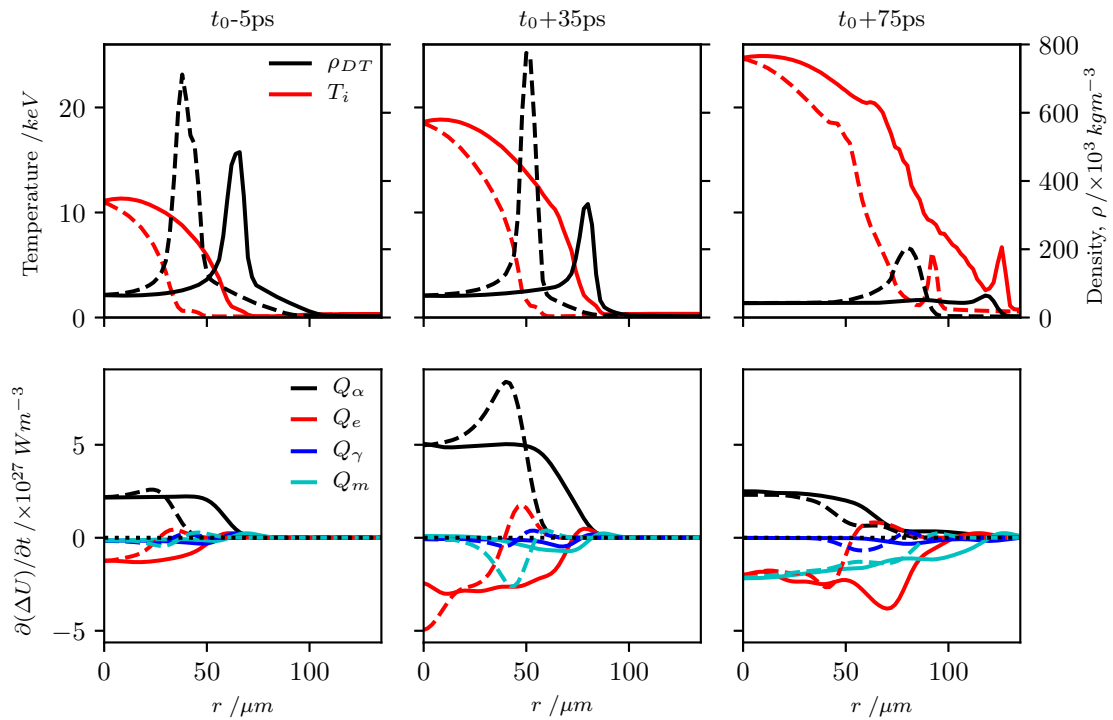
Figure 5.10: The time evolution of radial profiles directly into (solid) and away from (dashed) the perturbation bubble, showing (top) density (black) and ion temperature (red), and (bottom) the hotspot power balance contributions from alpha-heating, Q_α (black), electron thermal conduction, Q_e (red), radiation, Q_γ (blue) and mechanical work, Q_m (cyan).

rest of the shell against the central hotspot pressure. The resultant earlier re-expansion of this region from the hotspot pressure is what produces the bubble region. As the hotspot heats up and the pressure grows, both the forming bubble and the rest of the capsule start to re-expand, with the bubble re-expanding slightly faster and the hotspot material flowing towards the bubble.

Figure 5.10 shows the time evolution of the radial profiles along lines of sight directly into the bubble and away from the bubble. Profiles are shown for the ion temperature and density (above) and the hotspot power balance components (below) (similar to the time evolution of profiles for the single-spike in figure 5.2). With the perturbed bubble axis effectively expanding ahead of the rest of the shell, the temperature profiles seem to simply be stretched to match the larger extent along the bubble perturbation axis, while the shell peak density is reduced as well.

We note that this simulation exhibits characteristics of a robustly igniting capsule entering the propagating burn regime, in which the hotspot heats up significantly and begins to strongly ablate shell material, but the shell areal density is unable to contain rapidly growing pressure. Despite the inadequate confinement, the hotspot temperature continues to rise even after peak compression. This is due to the boosted velocity producing central hotspot temperatures of $T_i \approx 10\text{keV}$ by the time of peak compression, resulting in a very strong alpha-heating feedback loop. The subsequently shortened heating timescale of the hotspot allows the hotspot temperature to rise and burn the innermost region of the shell. However, the confinement time is still not long enough to allow burn propagation throughout the shell before the capsule disassembly.

Considering the power balance profiles, it can be seen that the increased distance between the shell and the centre of the hotspot along the perturbed axis removes the Bragg peak from the alpha-heating profile. This illustrates that the Bragg peak is primarily due to the proximity of the shell to the regions of strong fusion production, resulting in alpha-particles encountering the cold-dense material before depositing much energy in the hotspot, and depositing the majority of it into the shell. With both the perturbed and unperturbed axes sharing the same central conditions, the shell being further from the centre effectively extends the deposition region for alpha-particles to gradually deposit their energy in the hotspot, producing a gradual sigmoidal profile, similar to that seen in the self-heating regime.

The thermal conduction is also noteworthy along the perturbed axis in that there is very little deposition of heat flow along this radial direction. This is because the direction of the largest temperature gradient is no longer radial along the bubble's axis, but rather in directions transverse to the bubble's axis. Consequently, the gradient in the radial direction is reduced. Alternatively, this can be considered in terms of the increased divergence from the increased surface area of this region — the same amount of heat is flowing out through the region of solid angle, but is being directed out through a larger surface area, resulting in less energy deposited in a given part of the shell.

As in the previous section, we can again compare the properties between the hemisphere containing the bubble perturbation (the perturbed hemisphere) and the hemi-
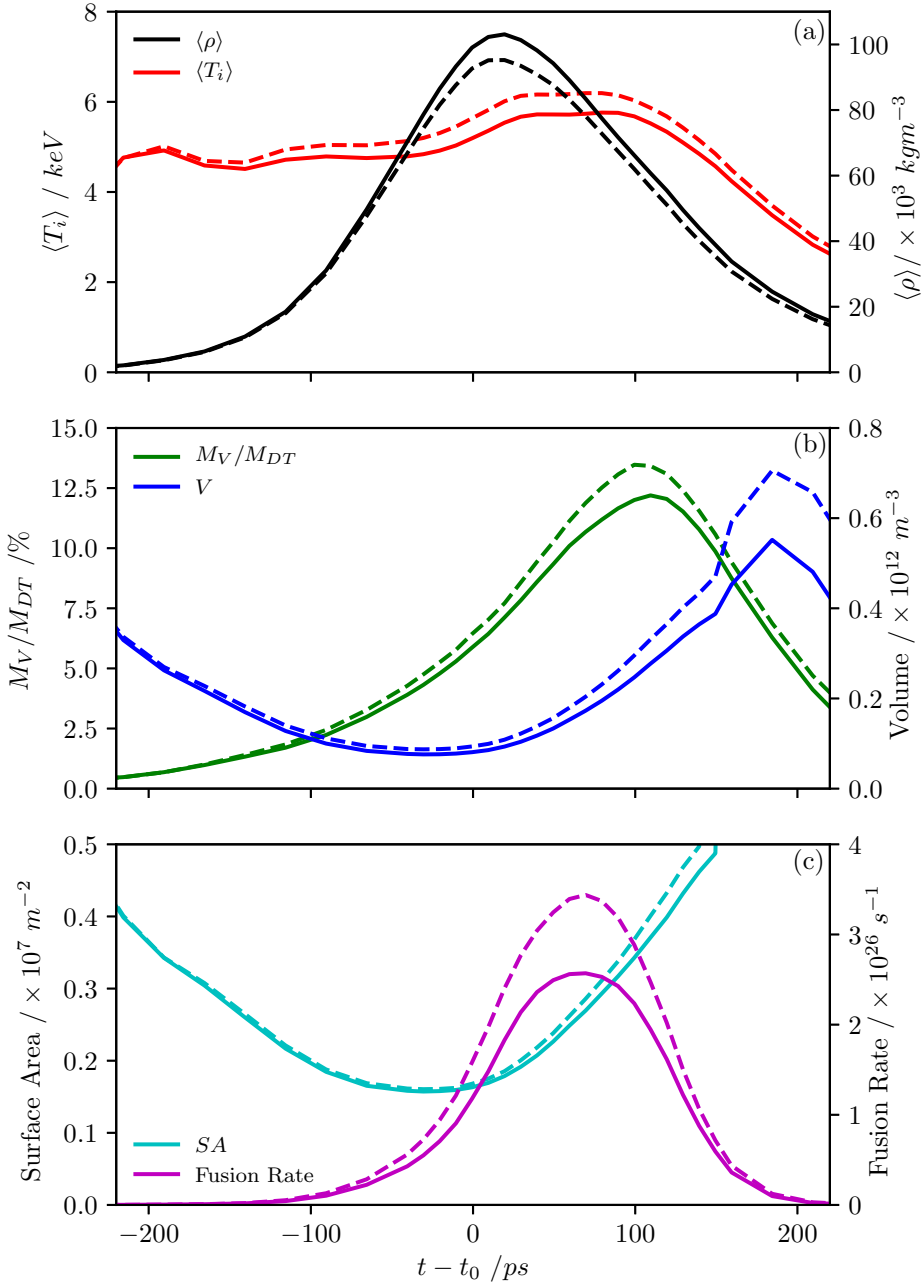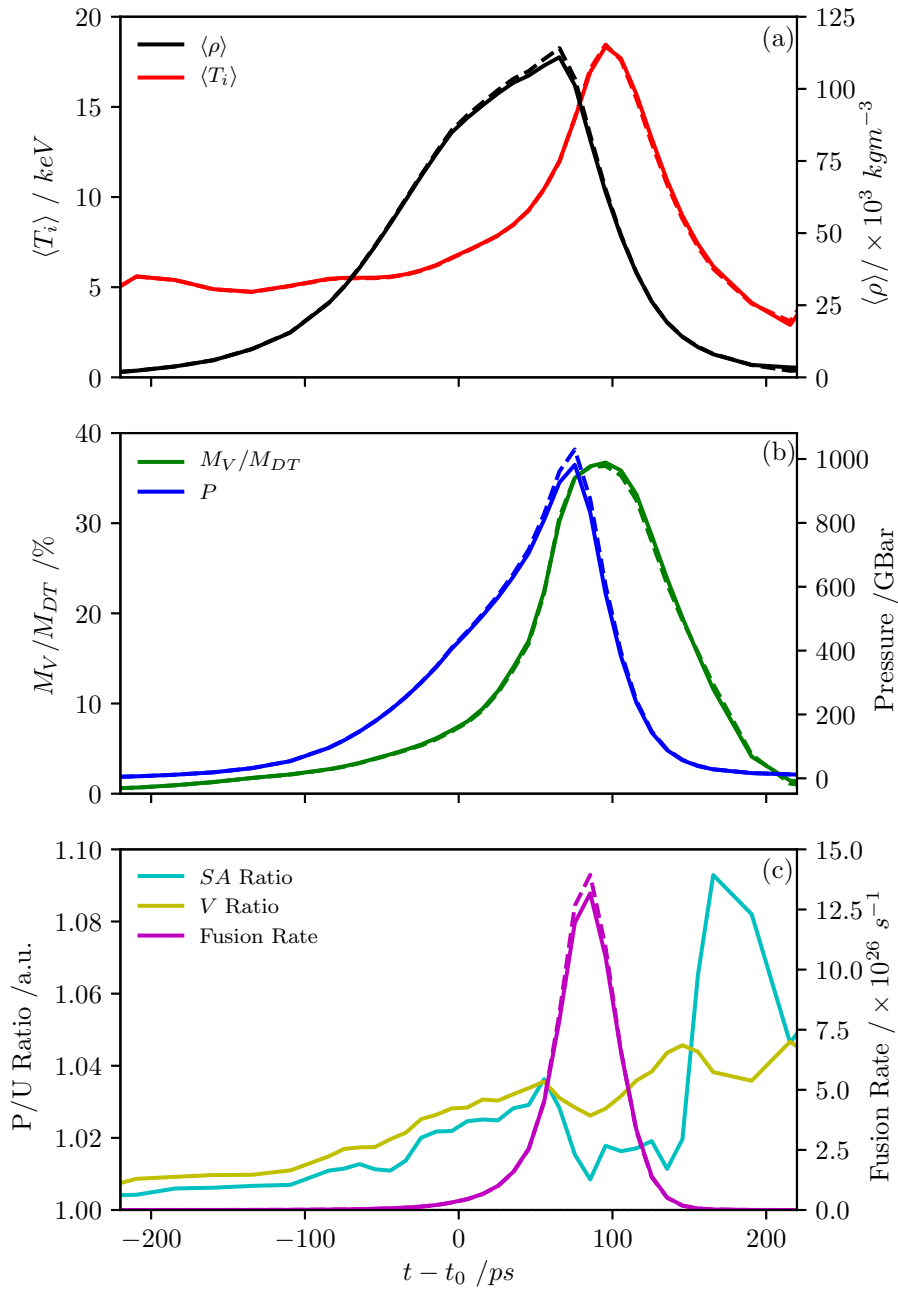
Figure 5.11: A comparison between the hemisphere containing the perturbation (solid) and the hemisphere without the bubble perturbation (dashed), showing: (a) hotspot burn-averaged density (black) and ion temperature (red); (b) pressure (blue) and mass as a fraction of the entire fuel (green); and (c) the DT fusion reaction rate (magenta) and the ratio of the surface area (cyan) and volume (yellow) between the perturbed and unperturbed hemispheres.
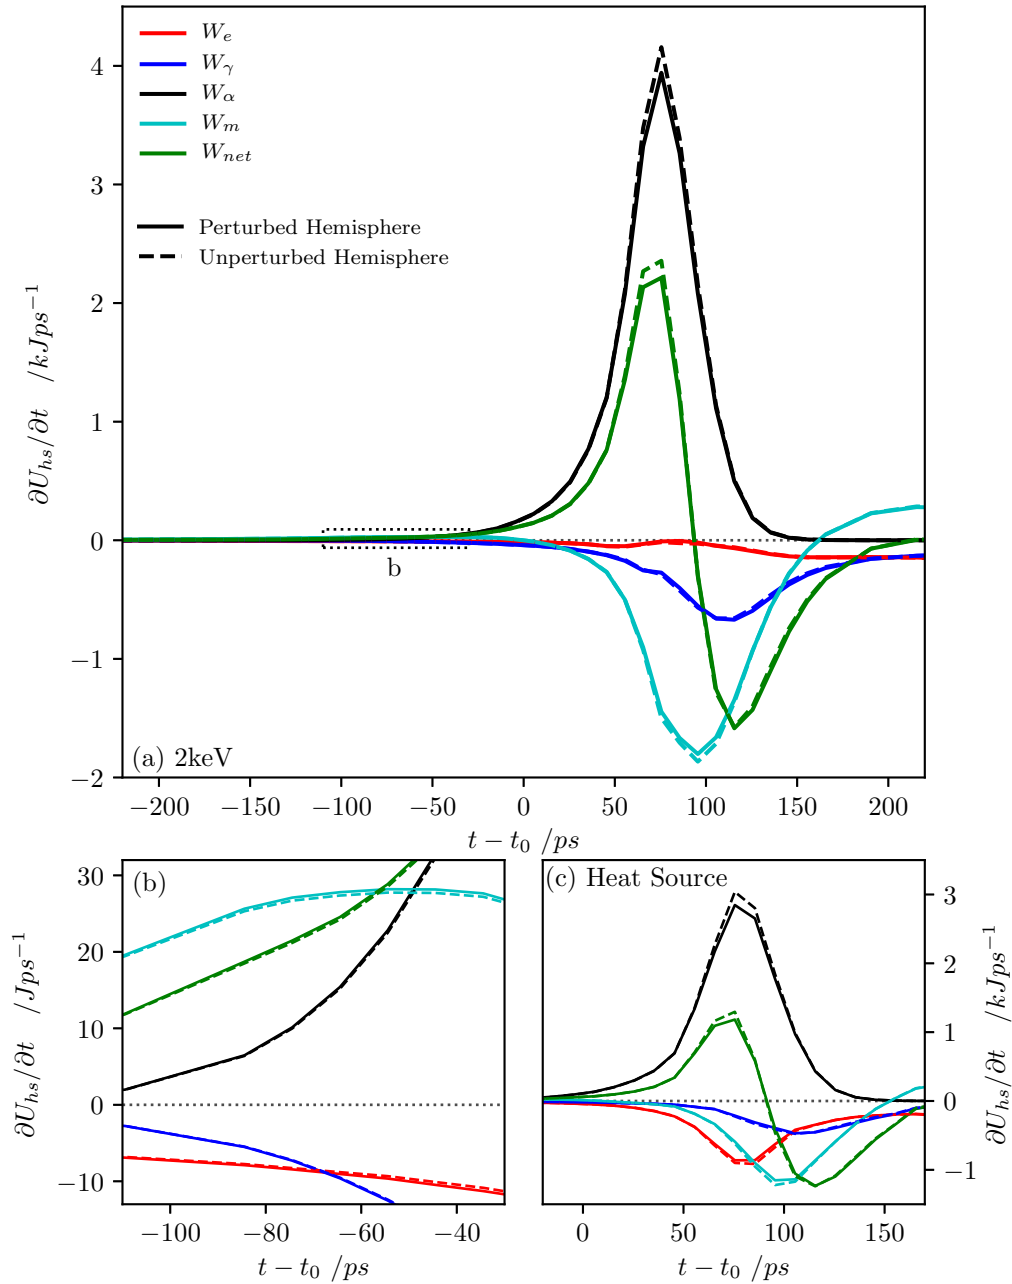
Figure 5.12:   (a) Volume-integrated power balance contributions for the 2keV hotspot
for the perturbed hemisphere (solid) and unperturbed hemisphere (dashed)
simulations, shown for alpha-heating, W_α (black), electron thermal conduc-
tion, W_e (red), radiation, W_γ (blue), mechanical work, W_m (cyan) and total,
W_{net} (green). (b) is zoomed in on the regions indicated by dashed box in
(a), showing the early-time variations during the deceleration phase, and (c)
shows the same quantities calculated for the heat-source hotspot.

sphere not containing the perturbation (the unperturbed hemisphere). Figure 5.11 compares the hotspot burn-averaged density, burn-averaged ion temperature, pressure, fractional mass, DT fusion reaction rate, surface area ratio and volume ratio between the two hemispheres. The ratios are calculated as the perturbed hemisphere quantity over the unperturbed hemisphere quantity. We note that the two hemispheres overall are very similar across all of the properties shown. The burn-averaged temperature is effectively identical across the entire time period shown. The burn-averaged density of the perturbed hemisphere is slightly below that of the unperturbed hemisphere after peak compression, corresponding to the earlier bubble re-expansion seen in figures 5.9 and 5.10. Both hemispheres contain a similar mass fraction, and therefore the larger volume of the perturbed case results in a lower density. The smaller increase in pressure is also linked to this, as the reduced confinement does not allow the pressure to properly build in this part of the hotspot — some of this pressure instead does work on the shell to expand the hotspot into this region of reduced confinement. The lower density material also therefore corresponds to less fusion occurring in the perturbed hemisphere.

Figure 5.12 shows the volume-integrated power balance contributions for the two hemispheres, calculated for the 2keV hotspot in (a) and the heat-source hotspot in (c). Figure 5.12(b) shows a zoom of (a) during the deceleration phase, as indicated by the dashed box in (a). The power balance comparison shows the slightly reduced levels of alpha-heating in the perturbed hemisphere, the peak of which occurs ~ 10ps before the peaks in fusion rate. The notably low thermal conduction rate is due to the $2keV$ contour being well beyond the region of deposition of the thermal conduction heat flow. The thermal conduction contribution to the power balance of the heat-source hotspot is significantly higher, although the other contributions are also significantly smaller due to the smaller volume of this type of hotspot.

In the exploration of this perturbation, we note that there was not enough areal density in the rest of the capsule to isolate the confinement loss to the developing hotspot bubble. While the bubble re-expands faster than the rest of the shell, the difference is not quite significant enough to produce loss of confinement through the bubble alone. The $S = 1.0$ capsule does not have the shell areal density to achieve this rapid localised loss of confinement due to low ρR, even though the increased velocity generated the pressures required to investigate this phenomena. The creation of a real "aneurysm" of a hotspot might require higher hotspot pressures or larger differences in shell areal density, such as with a higher average shell areal density or a larger amplitude perturbation. Note that both the pressure and the difference in shell areal density conditions could be achieved with higher scale factor capsule.

This simulation also highlights the difficulties in attempting to explore stronger burn regimes by simply increasing the implosion velocity of a particular design. Increasing the implosion velocity does indeed increase the total work done on the hotspot and generate higher temperatures in the hotspot. However, the correspondingly higher pressures are not also matched by better shell confinement. As a result, rather than

the hotspot burning significant proportions of the shell, the shell simply fails to confine the hotspot and is re-expanded before the burn wave can propagate through it. In other words, the hotspot converges faster, and then simply rebounds faster as well. Hydrodynamic scaling therefore combines the improved hotspot conditions with the enhanced confinement required to allow an igniting hotspot to propagate a burn wave through the shell.

5.4 Summary

In this chapter, we have described the methodology behind 3D simulations in this thesis. High resolution ($0.25\mu m$ or $0.5\mu m$) 1D and 2D spherical geometry simulations are used to accurately capture the physics of the ablation phase, complete with detailed radiation transport using 54 radiation groups to model the complex radiative dynamics of partially ionized carbon. This data is then remapped onto a lower resolution ($2\mu m$ or $3\mu m$) 3D Cartesian grid at peak radiation temperature, with any perturbations or hydrodynamic scaling (for chapter 6) applied at this reinitialisation point. A reduced 10-group radiation transport is used to capture the radiative dynamics of the DT in the hotspot during the stagnation and burn phases of the calculation. These 3D simulations are run without radiation drive, although this does not have a significant impact on the overall shell implosion velocity [130]. The velocity at reinitialisation is increased by $\lesssim 5\%$ in order to compensate for the reduced yield of lower resolution simulations.

A perturbation methodology used in both this chapter and chapter 6 is described, in which the perturbations are applied through a method developed by S. Taylor [91], using Layzer's approximate analytic treatment of the RTI [186] to perturb the velocity field of the perturbation. These velocity perturbations reproduce the expected behaviour of the RTI in the deceleration phase [91], and can be applied approximately uniformly across the surface of a sphere using the vertices of geodesic spheres, with different wavelengths dependent on the number of vertices. These perturbations can be pointed off-centre to induce non-spherical motion. In addition, coherent interference between multiple modes can be avoided through rotation of the vertices' co-ordinates, and the amplitudes can be randomised.

Using this method, a perturbation was applied to generate a single spike-like perturbation intruding into the hotspot in order to investigate its impact on the hotspot behaviour and capsule evolution. The spike effectively acts as a heatsink for the hotspot, drawing energy out of the hotspot and increasing the hotspot's energy losses. The thermal conduction losses increase due to the increased surface area, the increased temperature gradient, ∇T, around the spike, and an increase in the thermal conductivity due to the higher temperature at the location of the peak temperature gradient. Alpha-heating absorption increases into the spike, as the cold dense shell material is in closer proximity to the core of the hotspot; this increases the flux of the alpha-particles into the spike, and also increases the proportion of their energy deposited in the spike rather than the hotspot. The radiative losses are minimally affected by the spike. The overall

performance of the simulation is strongly impacted by these processes, both due to the enhanced losses after the hotspot formation, and a weaker alpha-heating feedback loop from the effect of poorer synchronisation of PdV work on the hotspot formation process.

A single bubble perturbation was applied to investigate confinement loss through a region of low areal density and therefore weak confinement. The bubble was found to have a small impact on the hotspot formation through the mistiming of the PdV compression during stagnation. However, unlike the spike perturbation, the bubble does not preferentially absorb heat from the hotspot and increase the heat flow out of the hotspot. The main effect occurs after the hotspot temperature and pressure have increased significantly; faster re-expansion into the region of weaker confinement than the rest of the shell acts to truncate burn by dropping the hotspot pressure.

From our simulations, we expect that a spike perturbation would be significantly more detrimental to the hotspot ignition process than a bubble perturbation. The significant increase in early time heat flow out of the hotspot will act to dampen the alpha-heating bootstrap behaviour, and potentially inhibit ignition. A bubble affects the hotspot confinement and can therefore truncate the burn propagation through confinement loss. However, it is likely to have a much smaller impact on the ignition of the hotspot due to the minimal effect on the heat flow out of the hotspot.

In these simulations, we have also examined the difficulties in attempting to reach the propagating burn regime through increases in implosion velocity alone. In addition to the increased susceptibility to perturbation growth, we find that although the higher implosion velocities result in more PdV work on the hotspot, higher temperatures and therefore a stronger alpha-heating bootstrap, the increased hotspot pressure is not matched by a corresponding increase in hotspot confinement. As such, the shell re-expands faster under the higher hotspot pressure before burn propagation can occur, although the hotspot is more likely to ignite robustly. Conversely, hydrodynamic scaling increases the shell areal density and therefore the confinement, while the hotspot heating is improved not by a higher initial temperature but by the longer duration for alpha-heating to occur, and a larger proportion of the alpha-heating being absorbed within the hotspot.

6 3D Perturbed Scaling

Recent experimental campaigns on NIF have been aimed at demonstrating the control and predictability of implosion designs through smaller scale designs and reduced laser energy [55, 188]. Although work continues towards reducing the magnitude of perturbations, the presence of perturbations will likely persist. Therefore, moving towards megajoule yield performance requires the scaling of these implosion designs [3] in order to improve performance while maintaining similar levels of symmetry control.

In this chapter, we hydrodynamically scale 3D perturbed hotspots in order to examine their behaviour and performance, particularly as they undergo ignition and burn. This is investigated for different forms of perturbation to consider the scaling of each perturbation scenario individually, rather than for a particular implosion design with a comprehensive set of perturbations, as in Clark *et al.* [3]. In the process, we also demonstrate the effects of alpha-heating on hydrodynamic scaling from Chapter 4 under 3D, perturbed conditions, and the resultant impact on the rest of the capsule.

Section 6.1 explores the hydrodynamic scaling for the Highfoot shot N130927, which was also the base implosion design used in Chapters 4 and 5, a full scale implosion shot at 1.8MJ [189]. In section 6.2, we explore this for a more recent HDC implosion N161023, originally shot at subscale $S = 0.8$ and 1.1MJ laser drive energy [55]. These HDC implosions have demonstrated improved symmetry control, and as a result have produced the highest yield experiments to date [9]. By exploring the scaling of two different designs, we can consider the impact of not only the different forms of perturbation, but also the yield degradation level, on the resultant performance scaling.

6.1 Highfoot Scaling

This section explores the impact of hydrodynamic scaling on perturbed implosions based on N130927. N130927 was a full-scale design with a 1.13mm outer radius capsule and a 5.75mm Au hohlraum shot at 1.8MJ using a 3-shock laser pulse [36, 189]. Despite being one of the better performing Highfoot implosions, N130927 still suffered from severe perturbations resulting in a $\sim 45\times$ yield degradation [2] compared to the expected symmetric implosion. We begin with an explanation of various aspects of the methodology behind this scaling study, before examining the impact of two highly degrading perturbations on the scaling of this design.
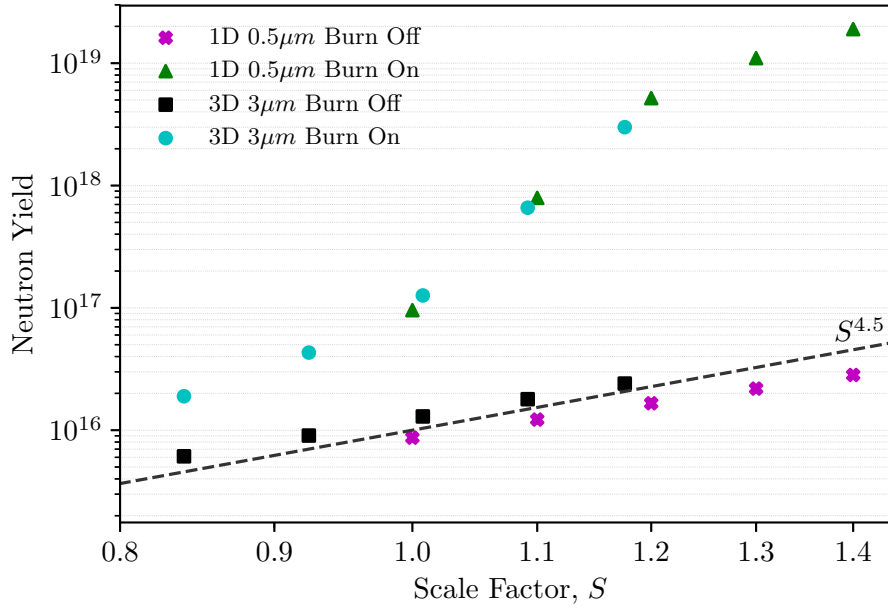
Figure 6.1: Comparison of neutron yield scalings for 3D $3\mu m$ and 1D $0.5\mu m$ simulations based on NIF highfoot shot N130927, for burn on (1D - green triangle, 3D - cyan circle) and burn off simulations (1D - magenta cross, 3D - black square). The 3D scales have been adjusted by a multiplier of 0.84 to fit the 1D curve. The dashed line indicates the expected burn off scaling of $S^{4.5}$ [128].

6.1.1 Methodology

As in Chapters 4 and 5, the simulations in this section are based on the same $0.25\mu m$ 1D simulation of N130927. Our 3D simulations here are run at $3\mu m$ resolution due to the dual computational constraints of simulating in 3D and the large number of simulations required in order to explore the scaling range. In addition, larger scale factors also require a corresponding increase in simulation domain size and runtime.

We note that the reduction in burn off yield of lower resolution 3D simulations also results in a diminished yield enhancement from alpha-heating, exacerbating the yield loss due to resolution. In this section, we can consider this as a reduction in the scale factor, S. Figure 6.1 shows the yields of 1D and 3D symmetric P_0 simulations at $0.5\mu m$ and $3\mu m$ respectively, for both burn on and burn off. Both the 1D and the 3D simulations were restarted from the same dataset at the same point in time, as described in section 5.1. The 3D simulations were run at scale factors S=1.0, 1.1, 1.2, 1.3 & 1.4, but are plotted with a multiplier of $0.84\times$ applied to the scale factor S in order to fit the curvature of the expected burn on 1D yield scaling. Although the 3D burn off yields do not match the 1D burn off yields with this multiplier, the hydrodynamic scaling is strongly affected by the alpha-heating and therefore we should not expect the same multiplier to fit both burn on and burn off. Indeed, using a multiplier of $0.94\times$ allows the 3D burn off yields to fit the 1D burn off yields.

The first perturbation scenario is based on the asymmetries in the radiation drive,
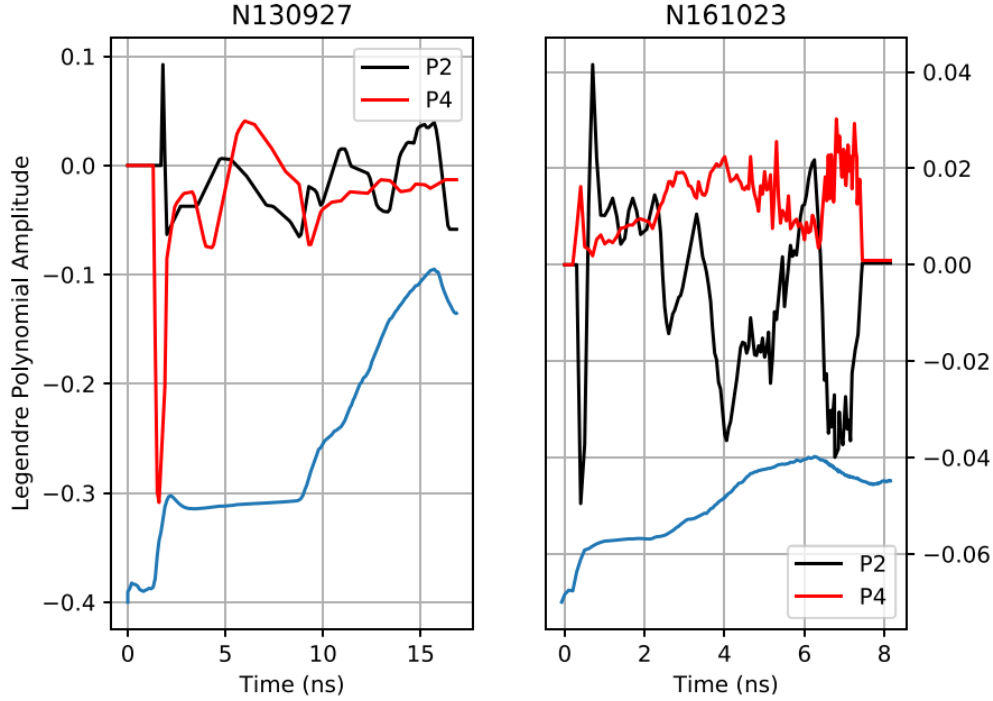
Figure 6.2: Time-dependent amplitudes of P_2 (black) and P_4 (red) Legendre modes applied in original 2D simulations of HF shot 130927 (left) and HDC shot 161023 (right), from which the corresponding RA scenarios in this chapter are reinitialised. The respective radiation drive pulses are also shown in arbitrary units (blue). Reproduced from [8] with the permissions of AIP Publishing.

labelled as the radiation asymmetry (RA) scenario. Data from 2D laser-hohlraum simulations [130] can be used to drive a 2D spherical geometry capsule simulation. The time-dependent amplitudes of the P_2 and P_4 Legendre modes used in this RA source simulation [8] are shown in figure 6.2 for both the HF scenario in this section, and the HDC scenario in the following section (6.2). This method results in the 2D source dataset at $S = 1.0$ inherently including effects such as mass redistribution [8], from which the 3D simulations are reinitialised at peak radiation temperature (i.e. $15.7ns$ here) and scaled. In doing so, we avoid the impact of the non-hydrodynamic scaling of the radiation drive, which would require adjustments to the ablator thickness [3] in order to obtain the correct optical depth, ablation front properties and implosion velocity.

The second perturbation scenario investigated is a short wavelength multi-mode (MM) perturbation. In this scenario, the 1D data is remapped onto the 3D Cartesian grid, and then perturbations applied at points distributed approximately uniformly around the surface of a sphere, as described in section 5.1. The vertex co-ordinates of different modes are rotated relative to one-another, and are each pointed off-centre (by a random distance up to $100\mu m$ in a random direction) to induce non-radial components to the perturbation. These perturbations are applied at an outer fuel radius of $394\mu m$ at $t = 15.7ns$, corresponding to the time of peak radiation temperature. The amplitudes are adjusted such that the simulation matches the yield of a comprehensive 3D simulation of the perturbations of N130927 [2] at the unadjusted scale factor S=1.0 ($\sim 3 \times 10^{15}$),
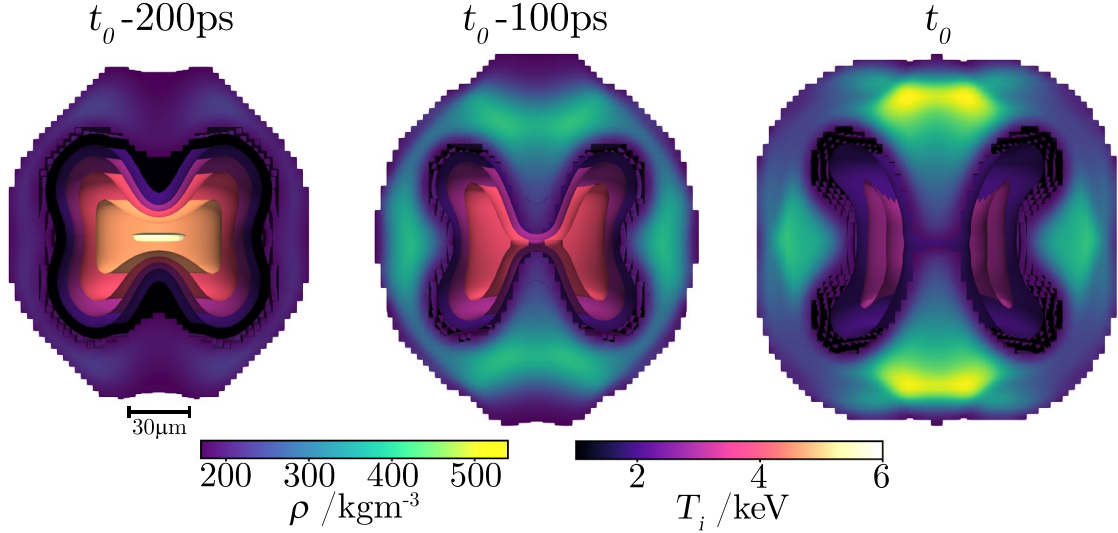
Figure 6.3: Hydrodynamic time evolution of the RA scenario showing the ion temperature
contours of the hotspot and the density of the shell, within $\sim \pm 100$ps of
bangtime. The central time-step $t_0 - 100$ps is at bangtime.

and are given below:

Wavelength, λ/R	Amplitude, $\delta v/\langle v_{imp} \rangle$
$1/\sqrt{12}$	0.15
$1/\sqrt{42}$	0.3
$1/\sqrt{162}$	0.36
$1/\sqrt{642}$	0.06
$1/\sqrt{2562}$	0.18

Table 6.1: Relative perturbation wavelengths, λ/R, and amplitudes, $\delta v/\langle v_{imp} \rangle$, for the
HF MM scenario.

6.1.2 Results and Discussion

Figures 6.3 and 6.4 show the hydrodynamic time evolution of the RA and MM scenarios,
showing the ion temperature contours in the hotspot, and the density of the shell around
bang-time, t_b, showing timesteps $t \sim t_b \pm 100$ps. These illustrate the shape and nature
of each of the perturbations; the multi-mode involves many spikes of various amplitudes
and wavelengths penetrating and cooling the hotspot, while the radiation asymmetry is
a P_2-P_4 swing in which the drive is initially stronger at the waist, and then stronger at
the pole. This drives perturbation spikes in from the poles of the capsule which penetrate
deeply into the hotspot and meet at the centre, effectively turning the hotspot (or what
remains of it) into a toroid. In both cases, bang-time is ~ 100ps before the time of peak
compression; for MM, $t_b = t_0 - 120$ps, while for RA, $t_b = t_0 - 100$ps.

Figure 6.5 shows the yield scaling of various 3D simulations, comparing the symmetric
burn on and off scenarios with RA and MM. These are plotted for the adjusted scale
factor, $\hat{S} = 0.84 \times S$ to fit the 1D scaling, as explained in section 6.1.1. Note that we will
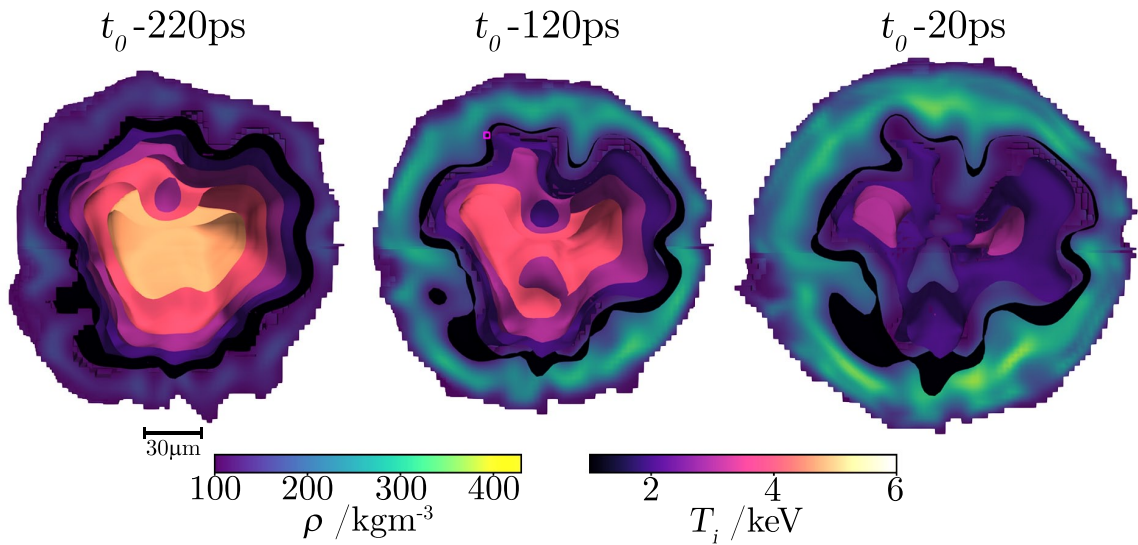
Figure 6.4: Hydrodynamic time evolution of the MM scenario showing the ion temperature contours of the hotspot and the density of the shell, within $\sim \pm 100$ps of bangtime. The central time-step $t_0 - 120$ps is 15ps before bang time.
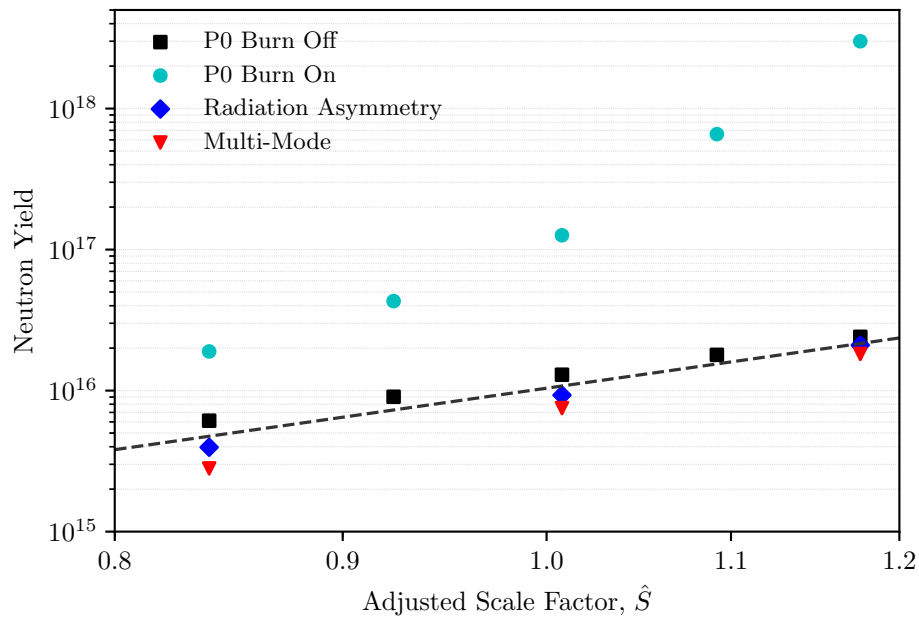
Figure 6.5: Neutron yield scaling of 3D simulations based on NIF highfoot shot N130927 for various implosion scenarios - symmetric burn off (black square), symmetric burn on (cyan circle), radiation asymmetries (blue diamond) and multi-mode A (red downwards triangle). The scale factor, S has been adjusted by a $0.84\times$ multiplier to give the adjusted scale factor, $\hat{S}$ which fits the 1D scaling. The dashed line indicates the theoretical burn-off scaling of $\hat{S}^{4.5}$.
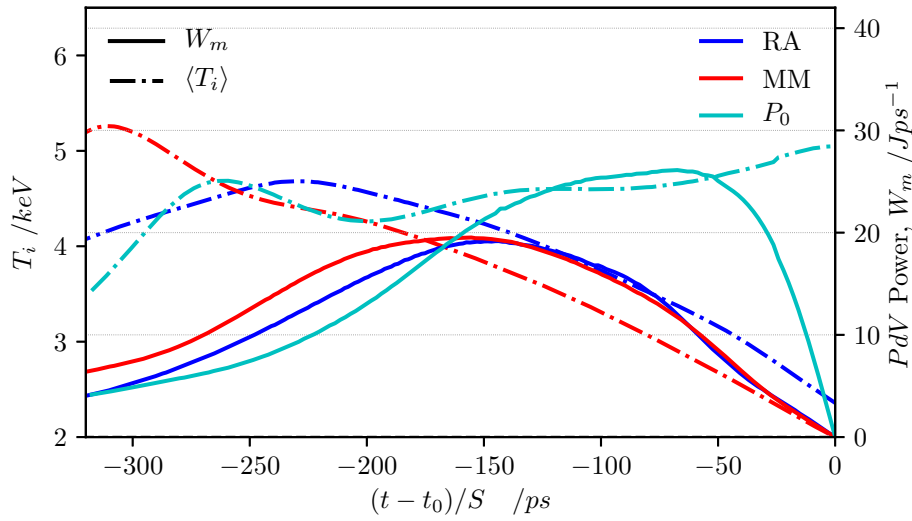
Figure 6.6:   The rate of PdV work done on the hotspot, W_m, (solid) and the hotspot
burn-averaged ion temperature, $\langle T_i \rangle$, (dash-dot) against time, for scenarios
P_0 (cyan), MM (red) and RA (blue) at scale factor S=0.8.

continue to refer to the scale factor of individual simulations by the unadjusted scale
factor S, for ease of notation (S=1.0, 1.2, 1.4 and $\hat{S}$=0.84, 1.008, 1.176). The symmetric
burn off case here scales as $\sim\hat{S}^{4.05}$, slightly below the $\hat{S}^{4.5}$ found in literature [128]. The
burn on scenario is beginning to curve upwards in the characteristic non-linearity of
ignition. Despite the lower yields, the perturbed burn on scenarios all scale faster in
yield than the burn off scenario. This is encouraging, as it indicates that alpha-heating
increases with scale and therefore improves yield scaling, even at low yields.

The radiation asymmetry, which was the dominant yield degradation mechanism for
N130927 [2], scales linearly as $\sim\hat{S}^{5.0}$, although the simulations perform worse than the
1D burn off scenario. The multi-mode scales faster at $\sim\hat{S}^{5.6}$, despite the initially lower
yield. This can be understood, in part, through the mechanical work delivered to the
hotspot. In figure 6.6, we examine the PdV work on the hotspot, W_m, and the hotspot's
corresponding (burn-averaged) ion temperature, $\langle T_i \rangle$. We note that the mechanical work
contribution, W_m, falls off at a very gradual gradient for both MM and RA scenarios,
compared to the unperturbed P_0 case.

Figures 6.7a and 6.7b quantify this, showing the scaled width of the PdV power
delivery, τ_m/S, where τ_m is defined as the FWHM of the W_m delivery:

$$\tau_m = t_b - t_a \tag{6.1}$$

where $t_{a,b}$ satisfy $W_m(t_{a,b}) = \frac{1}{2} \times \max(W_m)$ and $t_b > t_a$, and the total PdV work done
on the hotspot:

$$U_m = \int_{t_a}^{t_b} W_m(t)dt \tag{6.2}$$

These are plotted for MM, RA and P0 for S=1.0, 1.2 & 1.4 (for MM and RA only).
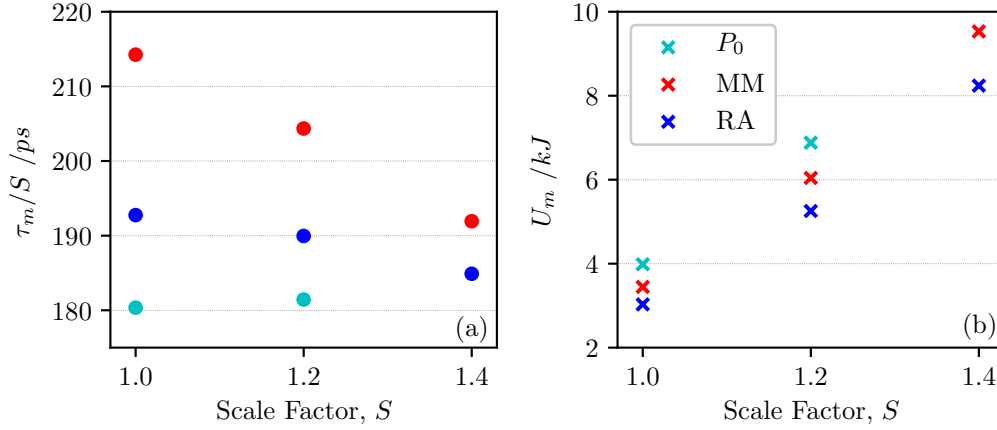
Figure 6.7: (a) Scaled PdV delivery width, τ_m/S (where τ_m is defined by the FWHM of W_m shown in 6.6) and (b) PdV work done on the hotspot, U_m (defined as the integral of W_m over the FWHM, as in equation 6.2) against scale factor for all scale factors, shown for P_0 (cyan), MM (red) and RA (blue) scenarios.

τ_m/S is largest for MM, coinciding with the earliest drop in $\langle T_i \rangle$, while the P_0 scenario has the shortest τ_m/S and largest boost to $\langle T_i \rangle$, in addition to the highest total work done, U_m. This indicates how the deceleration of the shell (and therefore the conversion of shell kinetic energy into hotspot thermal energy) is spread out over a longer time period in the perturbed scenarios, and therefore provides a much weaker boost to the hotspot temperature, also shown in figure 6.6. The slower compression allows more time for heat flow from the hotspot to the shell and therefore lower the temperature increase. Conversely, more synchronised compressions allow less time for heat flow and therefore the heating of the hotspot is more efficient.

Both the total work done, U_m, and the time-frame in which the work is done, τ_m/S, on the hotspot affect the quality of the PdV compression on the hotspot. Although MM does more work in total than RA, the time-frame of delivery is also longer, such that at $S=1.0$ the yield is lower ($\sim 3 \times 10^{15}$ compared to $\sim 4 \times 10^{15}$). However, U_m increases with $\hat{S}$ similarly across scenarios, and thus for a given increase in scale factor, the increase in U_m is higher for MM than for RA. In addition, τ_m/S decreases faster with $\hat{S}$ for MM, meaning that the improvement in the synchronisation of the PdV delivery is more significant for MM. This results in a larger increase in $\langle T_i \rangle$ with S for MM than for RA, which can be seen in figures 6.8a and 6.8b, and helps to explain, in part, the faster yield scaling of MM as compared to RA.

Figure 6.8 shows the burn-averaged ion temperature, $\langle T_i \rangle$, burn-averaged density, $\langle \rho \rangle$ and (volume-averaged) pressure, P, of the hotspot, as well as the burn history for the MM and RA scenarios at (unadjusted) scale factors $S=1.0$, 1.2 and 1.4. Figure 6.9 shows a side-by-side comparison of the net hotspot power balance and the alpha-heating contribution between the scale factors $S=1.0$, 1.2 and 1.4, with MM and RA shown for each scale

The burn-averaged ion temperatures, $\langle T_i \rangle$, are low in both perturbation scenarios,
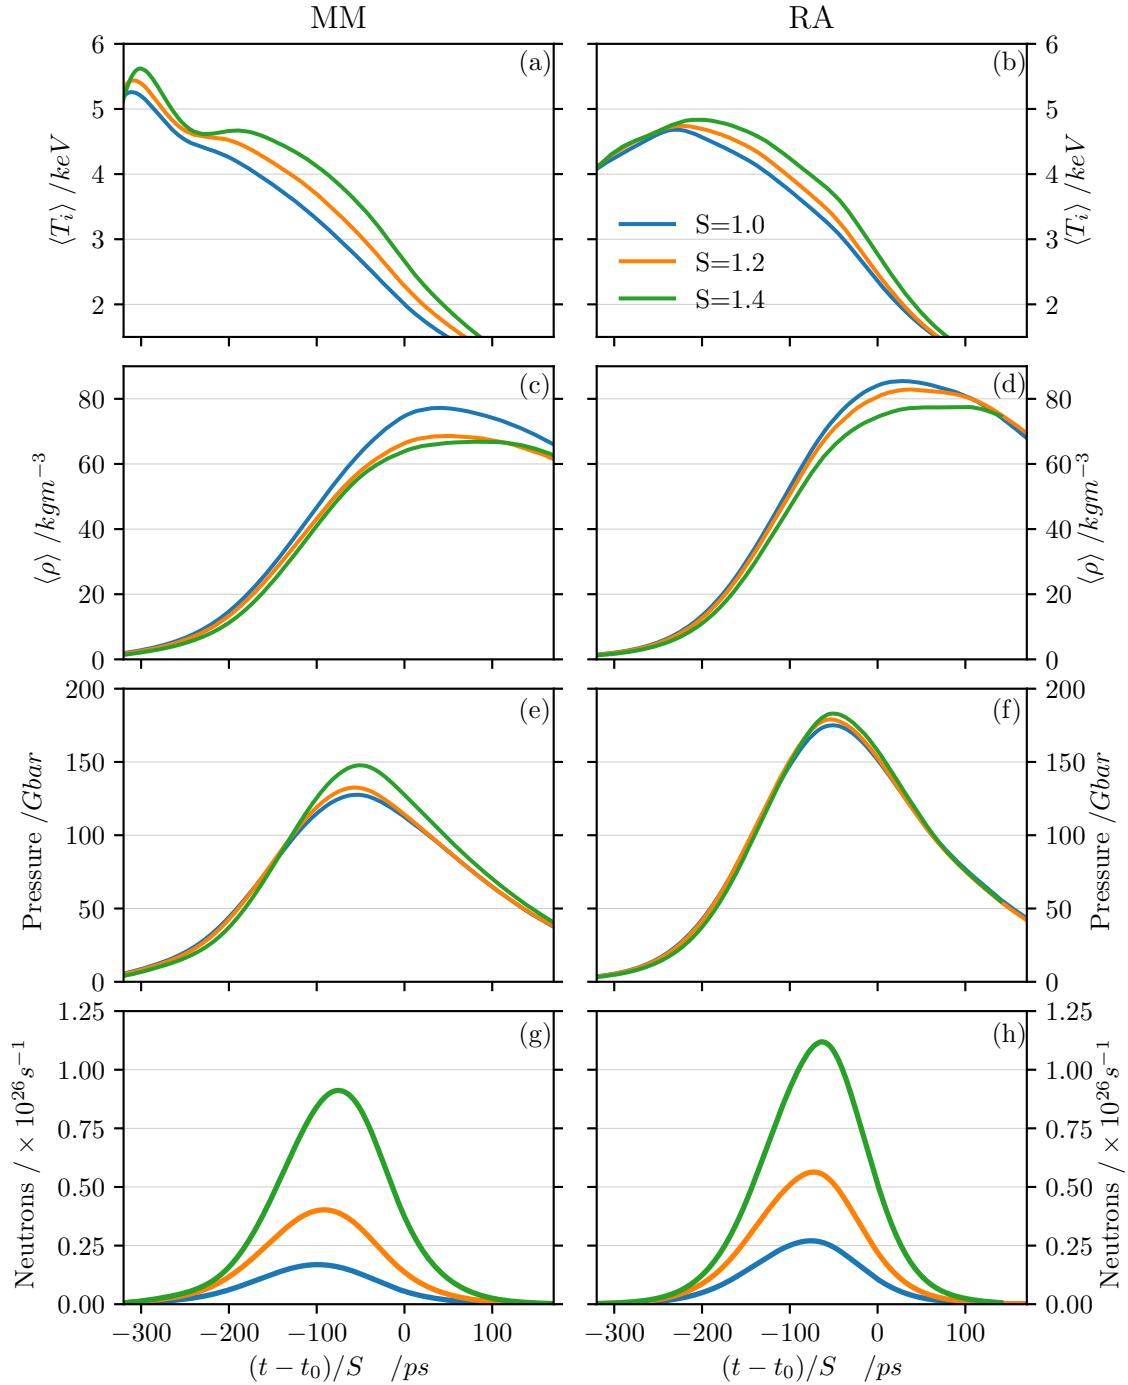
Figure 6.8: (a,b) Hotspot burn-averaged ion temperatures, $\langle T_i \rangle$; (c,d) hotspot burn-averaged density, $\langle \rho \rangle$; (e,f) hotspot pressure; and (g,h) burn history for MM (left) and RA (right) respectively, across scales: $S = 1.0$ (blue), $S = 1.2$ (yellow) and $S = 1.4$ (red). Scales are not adjusted by the 1D-3D multiplier here, but the native 3D scale. Times are normalised by the time of peak compression, t_0 and adjusted for scale factor S.
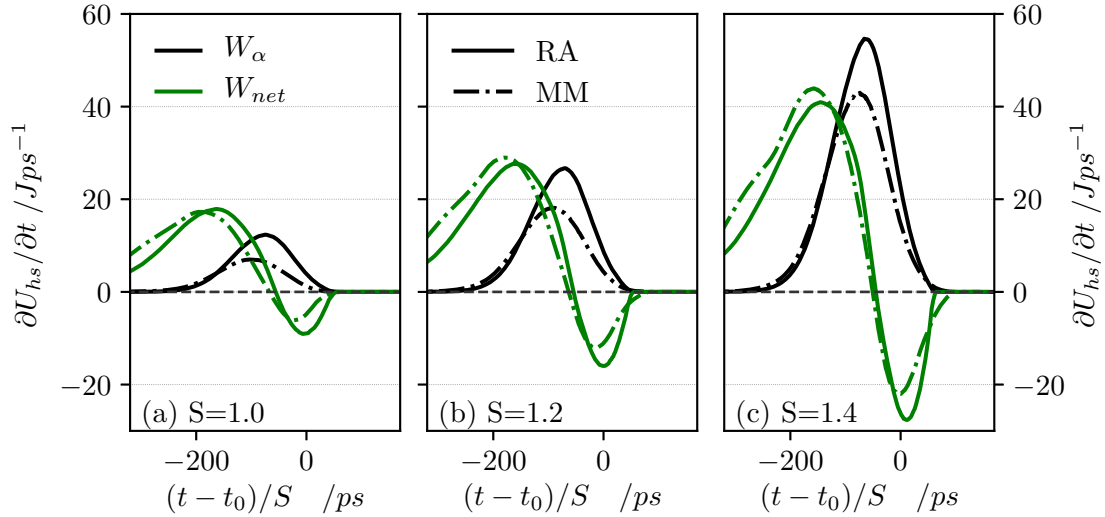
Figure 6.9: Comparison between RA (solid) and MM (dash dot) hotspot power balances for unadjusted scale factors (a) S=1.0, (b) S=1.2 and (c) S=1.4. The net power balance, W_{net} (green) is shown alongside the contributions from alpha-heating, W_α (black).

compared to the unperturbed scenario. For MM, the temperatures drop after the first shock flash (indicated by the first peak in $\langle T_i \rangle$), with a small boost in temperature coinciding with the peak of the PdV power delivery. For RA, this PdV peak coincides with the peak in $\langle T_i \rangle$, but in both cases the temperature drops beyond this time. The higher early time $\langle T_i \rangle$ with scale factor also corresponds to a decrease in $\langle \rho \rangle$, since the hotter material is harder to compress.

Neither scenario produces a significant level of alpha-heating — only for S=1.4 is the peak alpha-heating deposition rate, W_α in the hotspot larger than that of the mechanical work, W_m. Although not shown in figure 6.9 for clarity and simplicity, the mechanical work, W_m, is the most significant contributor to the early-time positive W_{net}. We would expect hydrodynamic scaling to improve the alpha-heating levels due to *a)* higher hotspot areal densities absorbing greater proportions of alpha-heating, *b)* a larger hotspot with more material producing more fusion reactions and *c)* lengthened implosion timescales allowing more fusion reactions to occur and produce alpha-heating before disassembly. However, the initial alpha-heating levels at S=1.0 are too weak in both perturbation scenarios for hydrodynamic scaling to produce a significant improvement in performance within this range of S. As such, the hydrodynamic scaling of these scenarios which struggle to progress even beyond the self-heating regime does not produce a significant yield enhancement. A significant proportion of the hotspot energy still stems from PdV work rather than alpha-heating, with total alpha-heating in the hotspot amounting to $\sim 35\%$ of the PdV work done on the hotspot for MM at $S = 1.2$ ($\hat{S} \approx 1.0$) and $\sim 55\%$ for RA.

MM scales faster than RA, but it can be seen that from the yield degradation levels of this implosion design, both scenarios still scale linearly and slowly. Achieving a high

yield from this starting yield would require a large increase in scale and laser energy. The yield scaling curve (figure 6.5) indicates that, increasing the scale factor to $\hat{S}$=1.2 from experimental yield levels might only result in a yield of $\sim 2 \times 10^{16}$. This would require laser energies of 3.1MJ to drive, and would only produce a yield which has already been achieved in the HDC campaign with a laser energy of 1.7MJ [9]. Although our perturbation scenarios are not a comprehensive set of realistic perturbations, this study remains informative about the prospects of scaling experiments from low yields, for both short- and long-wavelength perturbations. It is clear that increasing the scale factor on poorly performing existing designs will not be a viable way to reach ignition level yields. Instead, this indicates that reducing perturbation levels are necessary in order to achieve good performance, even if hydrodynamic scaling and larger lasers are to be used.

6.2 HDC Scaling

In this section, we explore the impact of hydrodynamic scaling of perturbed implosions based on the HDC [55] implosion design N161023 [3]. N161023 was a S=0.8 subscale implosion design consisting of a 0.91mm outer radius capsule with a 64μm thick layer of HDC, 19μm of which is doped with 0.22% by atomic weight of tungsten. This implosion was shot using 1.1MJ of laser energy in a 5.75mm depleted uranium (DU) hohlraum, and achieved good levels of symmetry control (within $\pm 5\mu m$ of a round implosion, even at a convergence ratio of 27), even at peak fuel velocities of 390km/s [55]. N170601 and N170827 were $S = 0.9$ versions of this design, shot at increased laser energies of 1.5MJ and 1.7MJ respectively to achieve yields of 1.7×10^{16} and 1.9×10^{16} [9], and therefore provide experimental data on the scaling performance of this implosion design. This was also used as the basis for HDC implosion scaling in a study extrapolating performance from current NIF experimental designs [3]. Therefore, N161023 is well-characterised, and provides a capsule design with a good initial performance from which to study the impact of hydrodynamic scaling on perturbations.

6.2.1 Methodology

As in section 6.1, the large numbers of simulations required in 3D present computational constraints which require us to use $3\mu m$ resolution simulations. However, in this section we instead apply a small multiplier to increase the velocity of our simulations at initialisation, rather than adjusting the results through a scale factor multiplier *post hoc*. This allows us to better explore the range of scale factors while keeping the simulation domain size as small as possible in order to minimise computational runtime.

In figure 6.10, we compare the yields of symmetric P_0 simulations run in 1D at 0.5μm resolution to those run in 3D at 2μm resolution with a velocity multiplier $s_v = 3\%$ and at 3μm resolution with $s_v = 4.8\%$. We can see that the use of a resolution-dependent multiplier, s_v, gives good agreement between the yields of the 2μm and 3μm
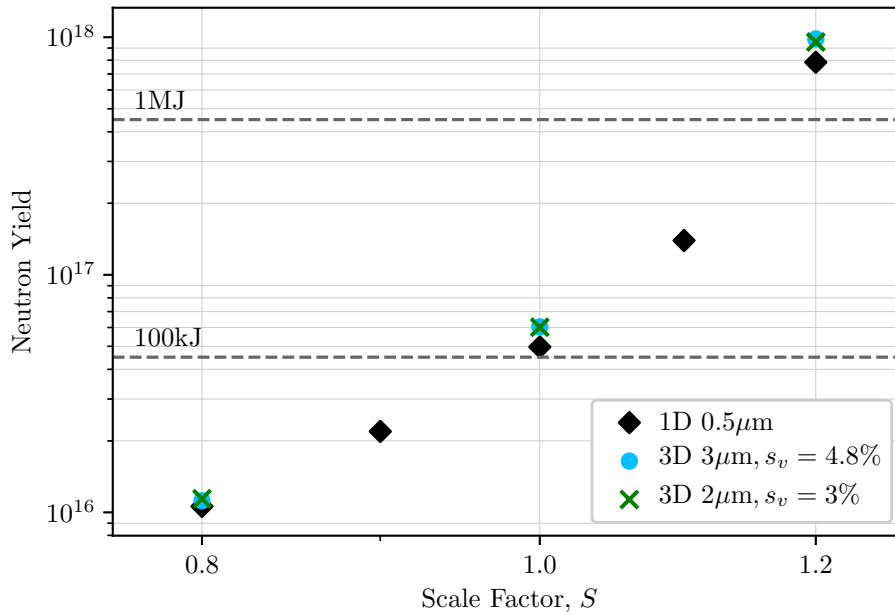
Figure 6.10: Yields as a function of scale factor on a log-log scale for symmetric P_0 simulations run at $0.5\mu m$ 1D (black diamond), $2\mu m$ 3D (red $\times$) and $3\mu m$ 3D (cyan circle). The 3D simulations have a velocity multiplier applied to increase the velocity by s_v; the $2\mu m$ 3D simulations are increased by $s_v = 3\%$ and the $3\mu m$ 3D by $s_v = 4.8\%$. $1MJ$ and $100kJ$ yield levels are marked in dashes.

3D simulations, and reasonable agreement between the 3D and 1D yields also.

This multiplier allows us not only to account for the effect of the reduced burn-off yield of lower resolutions inherently reducing the alpha-heating bootstrap effect, but also allows us to account for the lack of radiation drive. While this 'hydrodynamic decoupling' was previously verified [130], it was verified for normal scale capsules, rather than increased scale factor capsules which also have an increased implosion timescale. The lack of radiation drive allows the material at the outer edge of the capsule to trail behind and decompress, rather than remaining compressed against the core of the shell by the external drive. At normal scale factors, we do not expect this to have a significant effect on the confinement of the capsule as the timescales are too short for this to occur [130]. However, the lengthened timescales at increased scale factors could allow the outermost material to contribute to the confinement of the capsule, particularly since the increased scale factors will also increase the time duration of the burn pulse. Not accounting for this effectively equates to an increase in the coast-time — the time between the time at which the laser drive is shut off and the time of peak compression — which can affect the stagnation pressure and fusion yield [190].

As in sections 5.2, 5.3 and 6.1, we reinitialise our 3D simulations after peak radiation temperature (here at $t = 7.05ns$) from 1D and 2D spherical geometry calculations of the ablation phase, with the hydrodynamic scaling applied during this reinitialisation step. Similar to the Highfoot scaling study in section 6.1, we explore two types of

perturbation scenario: a short-wavelength multi-mode and a long-wavelength radiation asymmetry. All simulations in a particular scenario are initialised from the same dataset for reliability of comparison, as in the previous section.

The radiation asymmetry scenario (RA) is based on a 2D spherical geometry simulation with P_2 and P_4 Legendre modes applied in the radiation drive, as derived from 2D laser-hohlraum calculations [191]. The time-dependent amplitudes of these modes in the source HDC simulation were shown in the previous section (6.1)in figure 6.2 (right). The multi-mode scenario (MM) is initialised from a 1D simulation, and Rayleigh-Taylor spikes applied at multiple wavelengths approximately uniformly across the capsule surface (as described in sections 5.1.3 and 6.1), at an outer shell radius of $R = 309\mu m$. Here, the multi-mode amplitudes are tuned such that the S=0.8 multi-mode simulation yield matches the S=0.8 radiation asymmetry yield of $\sim 6.5 \times 10^{15}$, and are given in the following table: These simulations overpredict the yield compared to the experiment

Wavelength, λ/R	Amplitude, $\delta v/\langle v_{imp} \rangle$
$1/\sqrt{12}$	0.0
$1/\sqrt{42}$	0.4
$1/\sqrt{162}$	1.1
$1/\sqrt{642}$	1.7
$1/\sqrt{2562}$	1.6

Table 6.2: Relative perturbation wavelengths, λ/R, and amplitudes, $\delta v/\langle v_{imp} \rangle$, for the HDC MM scenario.

($\sim 4.5 \times 10^{15}$), which was subjected to multiple sources of perturbation in addition to the radiation asymmetries modelled (and matched by the multi-mode) here.

6.2.2 Differences in performance scaling

Figures 6.11 and 6.12 show the ion temperature contours in the hotspot and the density of the shell for timesteps $t \sim t_0 \pm 100$ps for the RA and MM HDC perturbation scenarios respectively. These show the shape and nature of the perturbations, similar to figures 6.3 and 6.4 in the previous section on Highfoot scaling. The multi-mode is significantly less distorted in shape than its Highfoot counterpart, which is a strongly contributing factor to the improvement in performance of these scenarios. The radiation asymmetry for N161023 is predominantly a P_4 perturbation, with rings of material constricting at polar angles $\theta = 45°, 135°$. This radiation asymmetry scenario is also weaker in amplitude than its Highfoot counterpart.

Figure 6.13 shows the yield scaling for 3D $3\mu m$ simulations of P0, MM and RA scenarios. Both MM and RA scenarios display signs of ignition in the upwards curvature on the log-log scale, with the yield increasing exponentially with scale factor. Both scenarios scale faster than the burn off 1D scaling of $S^{4.5}$, shown by the dotted line. Note that although at S=0.8 the MM and RA scenarios are matched in yield, the MM scenario yield increases faster with scale factor than the RA scenario.

t_0-90ps $\qquad$ t_0+10ps $\qquad$ t_0+110ps

45μm

200 $\qquad$ 400 $\qquad\qquad$ 2.5 $\quad$ 5 $\quad$ 7.5 $\quad$ 10

ρ /kgm⁻³ $\qquad\qquad$ T_i /keV
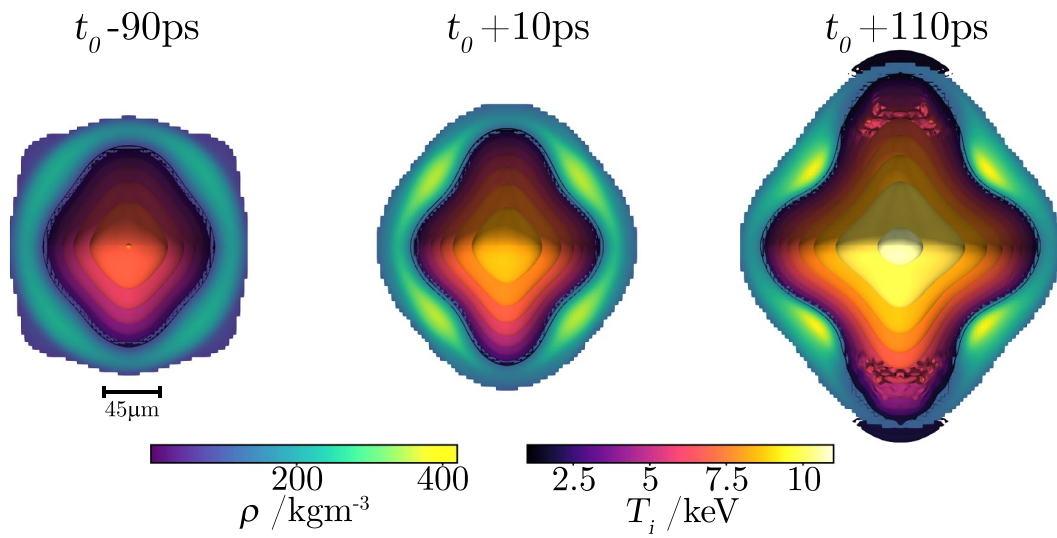
Figure 6.11: Time evolution of the RA scenario showing the ion temperature contours of the hotspot and the density of the shell, within $\sim \pm 100$ps of peak compression.

t_0-105ps $\qquad$ t_0-5ps $\qquad$ t_0+95ps

45μm

200 $\qquad$ 400 $\qquad\qquad$ 5 $\qquad$ 10

ρ /kgm⁻³ $\qquad\qquad$ T_i /keV
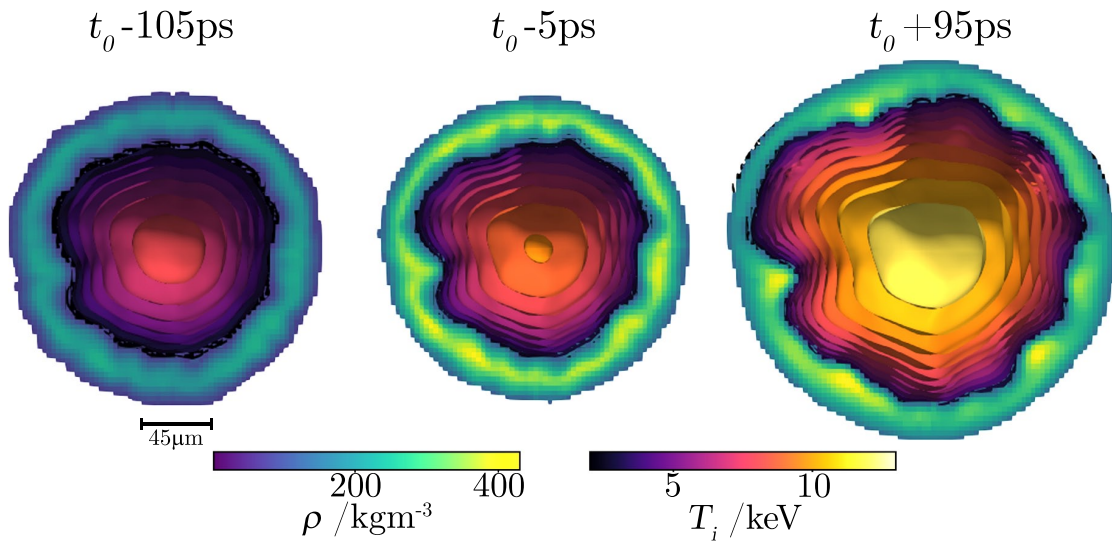
Figure 6.12: Time evolution of the MM scenario showing the ion temperature contours of the hotspot and the density of the shell, within $\sim \pm 100$ps of peak compression.
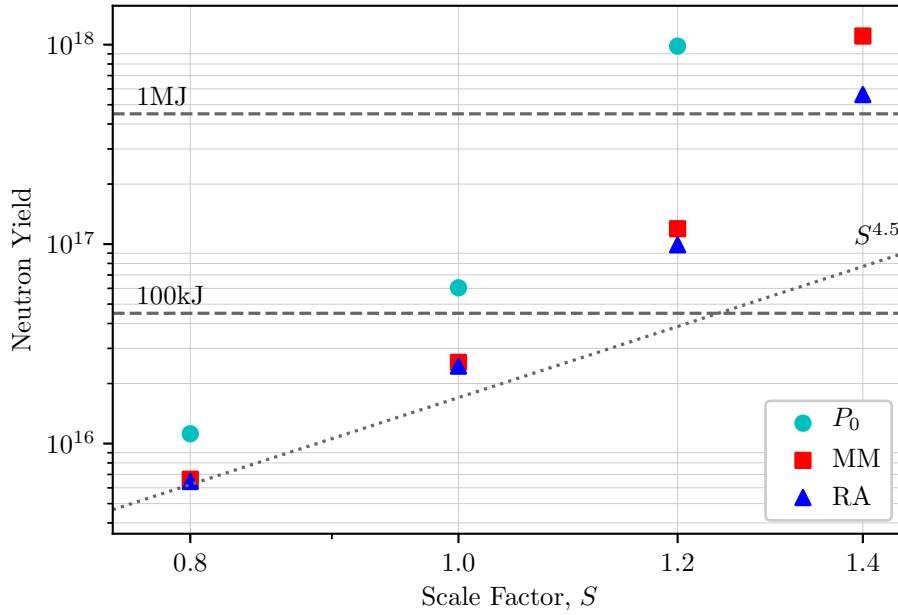
Figure 6.13: Yields as a function of scale factor shown on a log-log scale for 3D simulations of the symmetric P_0 (cyan circle), the multi-mode (red square) and the radiation asymmetry (blue triangle) scenarios. These simulations are run at $3\mu m$ resolution with a velocity multiplier of $s_v = 4.8\%$. $1MJ$ and $100kJ$ yield levels are marked in dashes, and the 1D, no-alpha yield scaling ($S^{4.5}$) is shown as the dotted line.
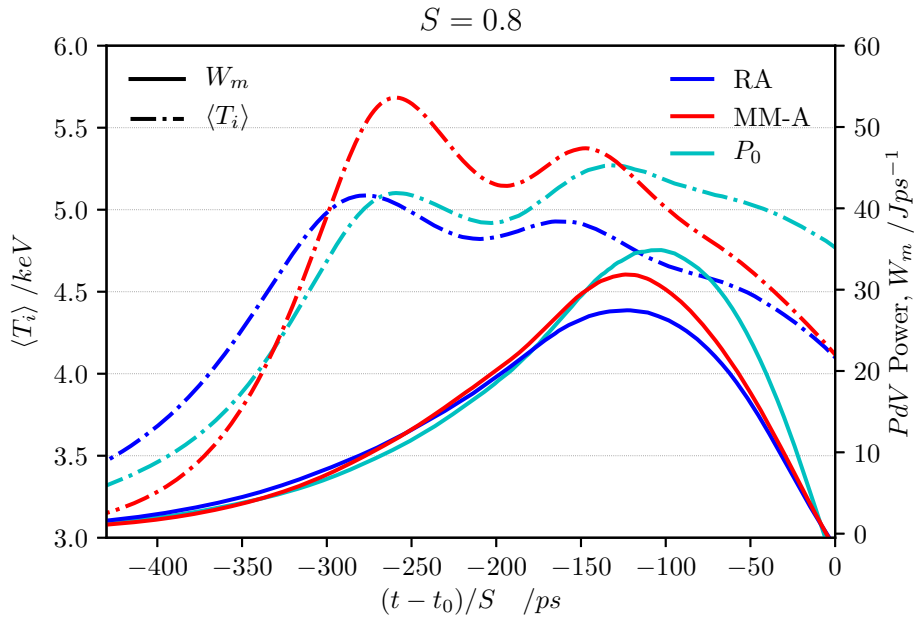
Figure 6.14: The rate of PdV work done on the hotspot, W_m (solid) and the hotspot burn-averaged ion temperature, $\langle T_i \rangle$ (dash-dot), against time for scenarios P_0 (cyan), MM (red) and RA (blue) scenarios at $S=0.8$. Time is shown relative to peak compression, $t - t_0$, and adjusted for scale factor, S.
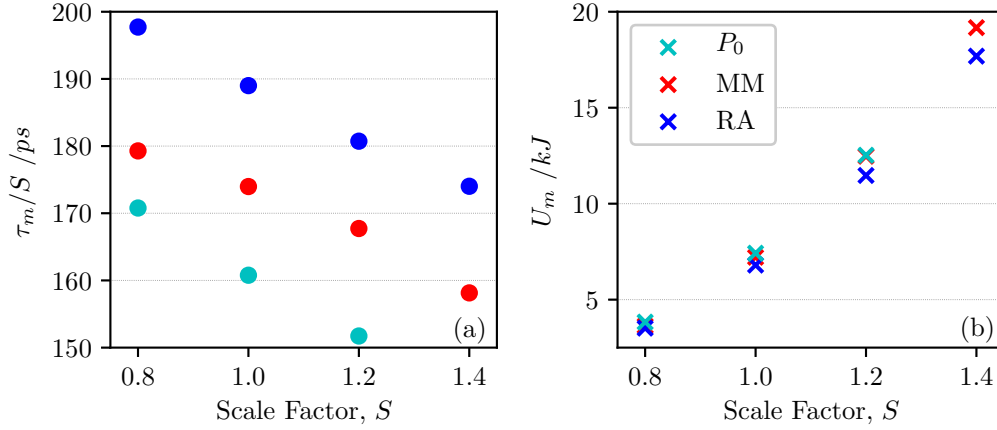
Figure 6.15: (a) Scaled PdV power delivery width, τ_m/S and (b) PdV work done on the hotspot, U_m, against scale factor for HDC scaling P_0 (cyan), MM (red) and RA (blue).

To understand the reason for this, we again examine the rate of PdV work done on the hotspot and its relationship to the burn-averaged ion temperature, $\langle T_i \rangle$ in figure 6.14. Compared to the equivalent figure 6.6 in section 6.1, we note that the PdV delivery of both MM and RA are qualitatively more similar to that shown for P_0 for HDC than for the Highfoot scalings. Between the HDC scenarios, it is clear that the MM scenario is more 1D-like than the RA scenario, while the most PdV work is done on the hotspot for the P_0 case.

In the symmetric case, all regions of the shell have the same implosion velocity, and will therefore stagnate simultaneously. However, in the perturbed case, velocity variations throughout the shell will cause different regions to stagnate at different times, spreading the PdV work over a longer time period. The effect of this is more prominent for RA than for MM. This may be due to the mass redistribution [8] from the longer wavelength perturbation, which results in larger variations in shell momentum than for the smaller short wavelength MM spikes, and therefore affects the time synchronisation and efficiency of the shell stagnation.

Figure 6.14 illustrates the relationship between the time synchronisation of the PdV work and the compression heating of the hotspot. Although initially the PdV power for MM is higher than for P_0, the P_0 power peaks higher, which coincides with a larger increase in the hotspot burn-averaged temperature. Both $\langle T_i \rangle$ and W_m are generally lower for RA, with the lower peak W_m coinciding with a lower increase in $\langle T_i \rangle$. More synchronised PdV stagnation allows less time for the hotspot to lose heat to the shell during the compression, and thereby results in better heating of the hotspot.

Quantifying this, the scaled width of the PdV delivery, τ_m/S (equation 6.1) and the total PdV work done, U_m (equation 6.2) are shown for all scales and scenarios in figures 6.15a and 6.15b respectively. Here, we can see that the more synchronous (i.e. shorter τ_m/S) PdV delivery extends to larger scale factors, with τ_m lowest for the P_0 scenarios, and larger for RA than MM. Similarly, the total PdV work done on the hotspot, U_m
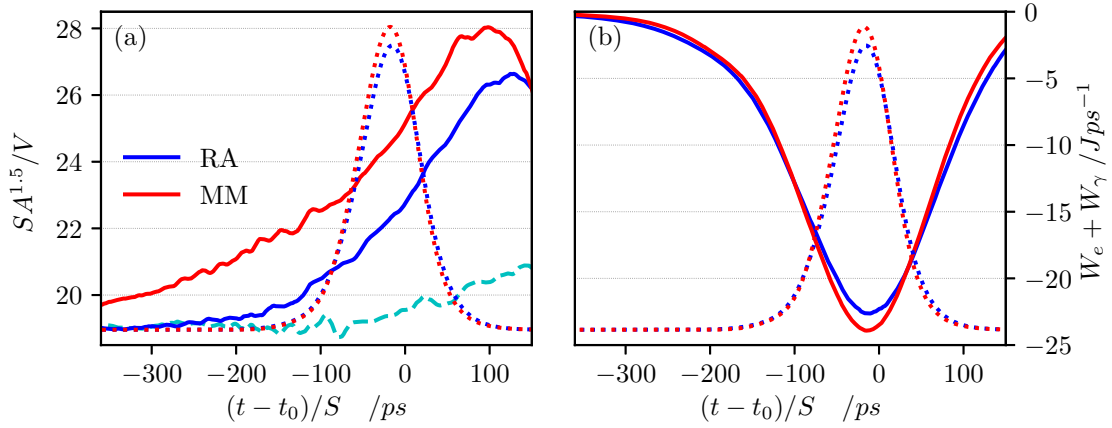
Figure 6.16: (a) The dimensionless surface area to volume ratio, $(SA)^{1.5}/V$ and (b) the combined thermal conduction and radiation power losses from the hotspot for MM (red) and RA (blue) scenarios at $S{=}0.8$. The dotted lines show the normalised neutron production rate to indicate the timing of burn. In (a), $(SA)^{1.5}/V$ is also plotted for the P_0 scenario (cyan dashed) as a point of reference.

is largest for P_0 scenarios and lowest for RA. The higher PdV power delivery results in better heating and higher temperatures of the hotspot, a trend that continues to higher scale factors.

However, in spite of more PdV work being done on the hotspot in the MM scenario, the overall performance still matches that of RA at $S{=}0.8$. Figure 6.16a shows the dimensionless surface area to volume ratio of the hotspot, $(SA)^{1.5}/V$ for all three scenarios at $S{=}0.8$, with normalised neutron production rate of MM and RA shown in the dotted line. The combined rate of energy loss from the hotspot due to thermal conduction, W_e, and radiation W_γ for MM and RA at $S{=}0.8$ is shown in figure 6.16b. Figure 6.17 shows the net hotspot power balance, W_{net}, for MM and RA across scale factors $S{=}0.8$, 1.0, 1.2 & 1.4, in addition to the contribution from alpha-heating, W_α.

For a perfect sphere, $(SA)^{1.5}/V$ is constant $(= 6\sqrt{\pi})$. The P_0 scenario remains roughly constant (albeit $\neq 6\sqrt{\pi}$ due to spatial discretisation). The MM has a consistently higher ratio than for either of RA or P_0, indicating that it is the least spherical in shape of the scenarios. The increased temperature, in addition to the larger $(SA)^{1.5}/V$ ratio, contribute to an increased thermal conduction loss rate ($\sim 10\%$) for MM compared to RA.

Although the PdV work produces an initially higher $\langle T_i \rangle$ for MM, the MM also cools faster such that over the duration of burn, the hotspot temperature and density are only slightly higher for MM than RA. Both scenarios at $S{=}0.8$ produce similarly small levels of alpha-heating, but slightly higher for MM due to the better PdV delivery and higher temperature. However, the small differences in $\langle T_i \rangle$ and $\langle \rho \rangle$ also cause a higher radiation loss rate ($\sim 5\%$) for MM. The higher pressure also results in a faster re-expansion, causing the hotspot to disassemble faster and reducing the burn duration.
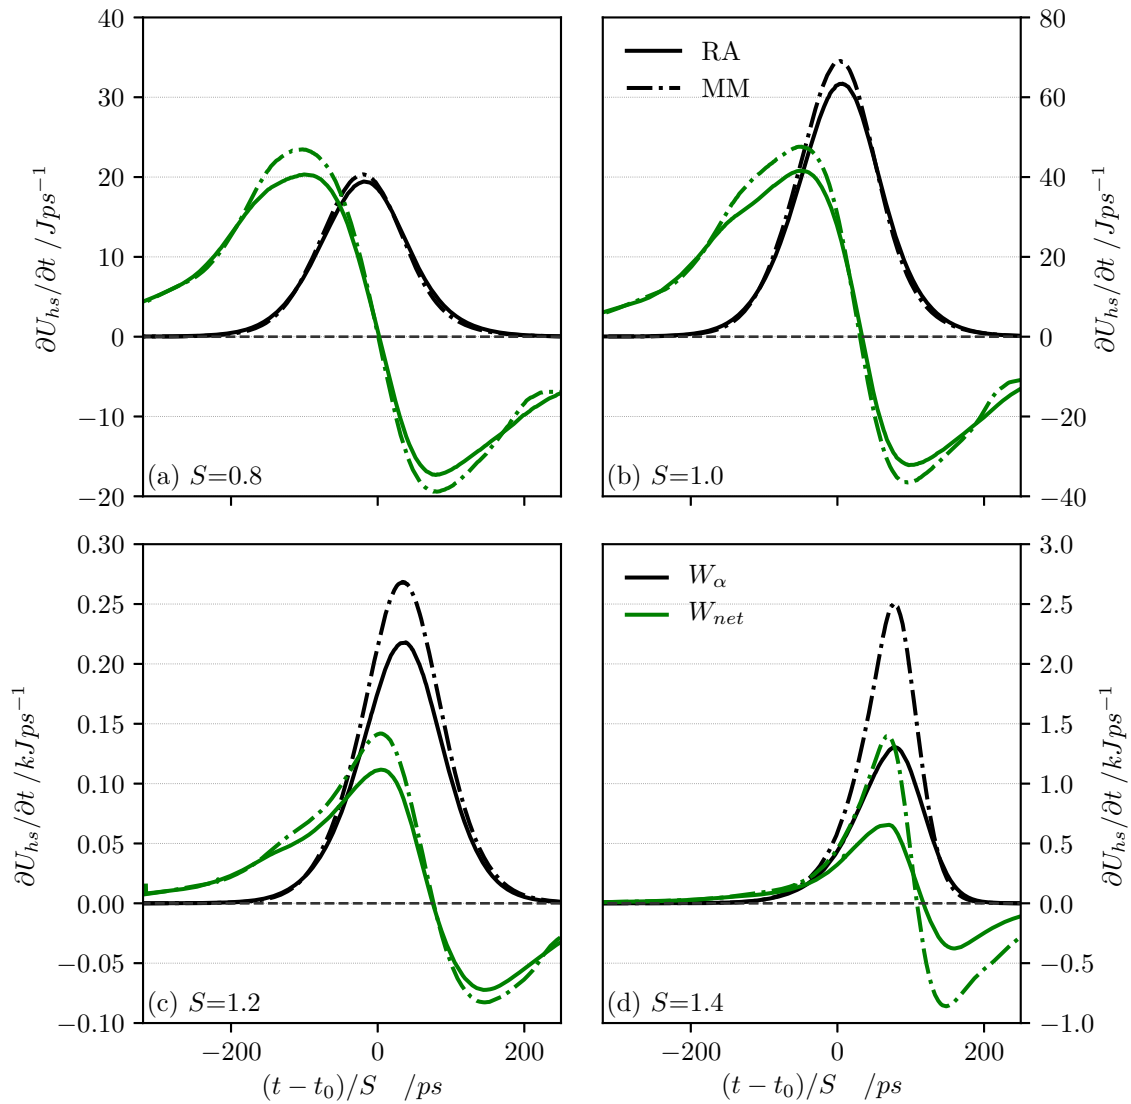
Figure 6.17: Volume-integrated hotspot power balance contributions for (a) S=0.8, (b) S=1.0, (c) S=1.2 and (d) S=1.4 for the MM (dash-dot) and RA (solid) scenarios, showing the net power balance, W_{net} (green) and the alpha-heating contribution, W_α (black).

Overall, this results in a similar performance for MM to RA at $S=0.8$.

As S increases, alpha-heating becomes increasingly dominant in the overall power balance. The larger radius and volume of the hotspot result in more mass contained within the hotspot, while the density and temperature profiles (and therefore the fusion reaction profile also) remain self-similar. This therefore results in more fusion reactions overall within the hotspot, and more alpha-particles to deposit and heat the hotspot at the start of the alpha-heating bootstrap. In addition, the larger hotspot areal density also retains more of the alpha-particle energy, and the longer timescale of higher scale factor implosions allow more time for fusion reactions to occur.

The more synchronous PdV compression in the MM scenarios (in which more work is done on the hotspot in a shorter time period) produce higher hotspot temperatures and densities at peak compression, and therefore result in stronger ignition in MM than in RA due to the nature of the alpha-heating bootstrap process. As in the Highfoot scaling in section 6.1, the short wavelength perturbation (MM) produces better yield scaling than the long wavelength perturbation (RA).

6.2.3 Effects of scaling on perturbed hotspots

Figures 6.18 and 6.19 show grids of plots of the MM and RA hotspots as a function of scale factor (across the grid) and time after peak compression (down the grid), displaying 2D slices of density (left half) and ion temperature (right half). The physical scale is normalised by scale factor in order to better compare the hydrodynamics. The side-by-side comparison of the time evolution of the different scale factors (and therefore different burn regimes) demonstrate not only the effects of alpha-heating on hydrodynamic scaling shown in section 4.3, but also demonstrate effects particular to perturbed 3D scenarios.

Figure 6.20 shows, for scenarios (MM,RA) respectively, the $2keV$ hotspot's: (a,b) burn-averaged ion temperature, $\langle T_i \rangle$; (c,d) burn-averaged density, $\langle \rho \rangle$; (e,f) mass relative to total DT fuel mass, M_{hs}/M_{DT}; (g,h) pressure; and (i,j) dimensionless surface area to volume ratio, $(SA)^{1.5}/V$.

The alpha-heating regime of each simulation is indicated by $\langle T_i \rangle$, $\langle \rho \rangle$ and M_{hs}/M_{DT} in figure 6.20. $S=1.4$ MM demonstrates propagating burn in a perturbed hotspot, shown by the sustained increase in $\langle T_i \rangle$ beyond peak compression in addition to significant ablation of shell material; this is indicated by the large increase in M_{hs}/M_{DT} and the broadened peak in $\langle \rho \rangle$. Similar features are seen for $S=1.4$ RA but less prominent, suggesting a weaker level of burn; the increase in $\langle T_i \rangle$ is slower and smaller in magnitude, and less material is ablated, such that the increase in hotspot mass is $\sim 20\%$ lower and the peak broadening in the time evolution of $\langle \rho \rangle$ is less pronounced. Both $S=1.2$ MM and RA show robust ignition, with $\langle T_i \rangle$ remaining relatively stable around peak compression, and a modest increase in hotspot mass. All simulations at $S=0.8$ & 1.0 show falling temperatures at peak compression, and no increase in hotspot mass after peak compression.
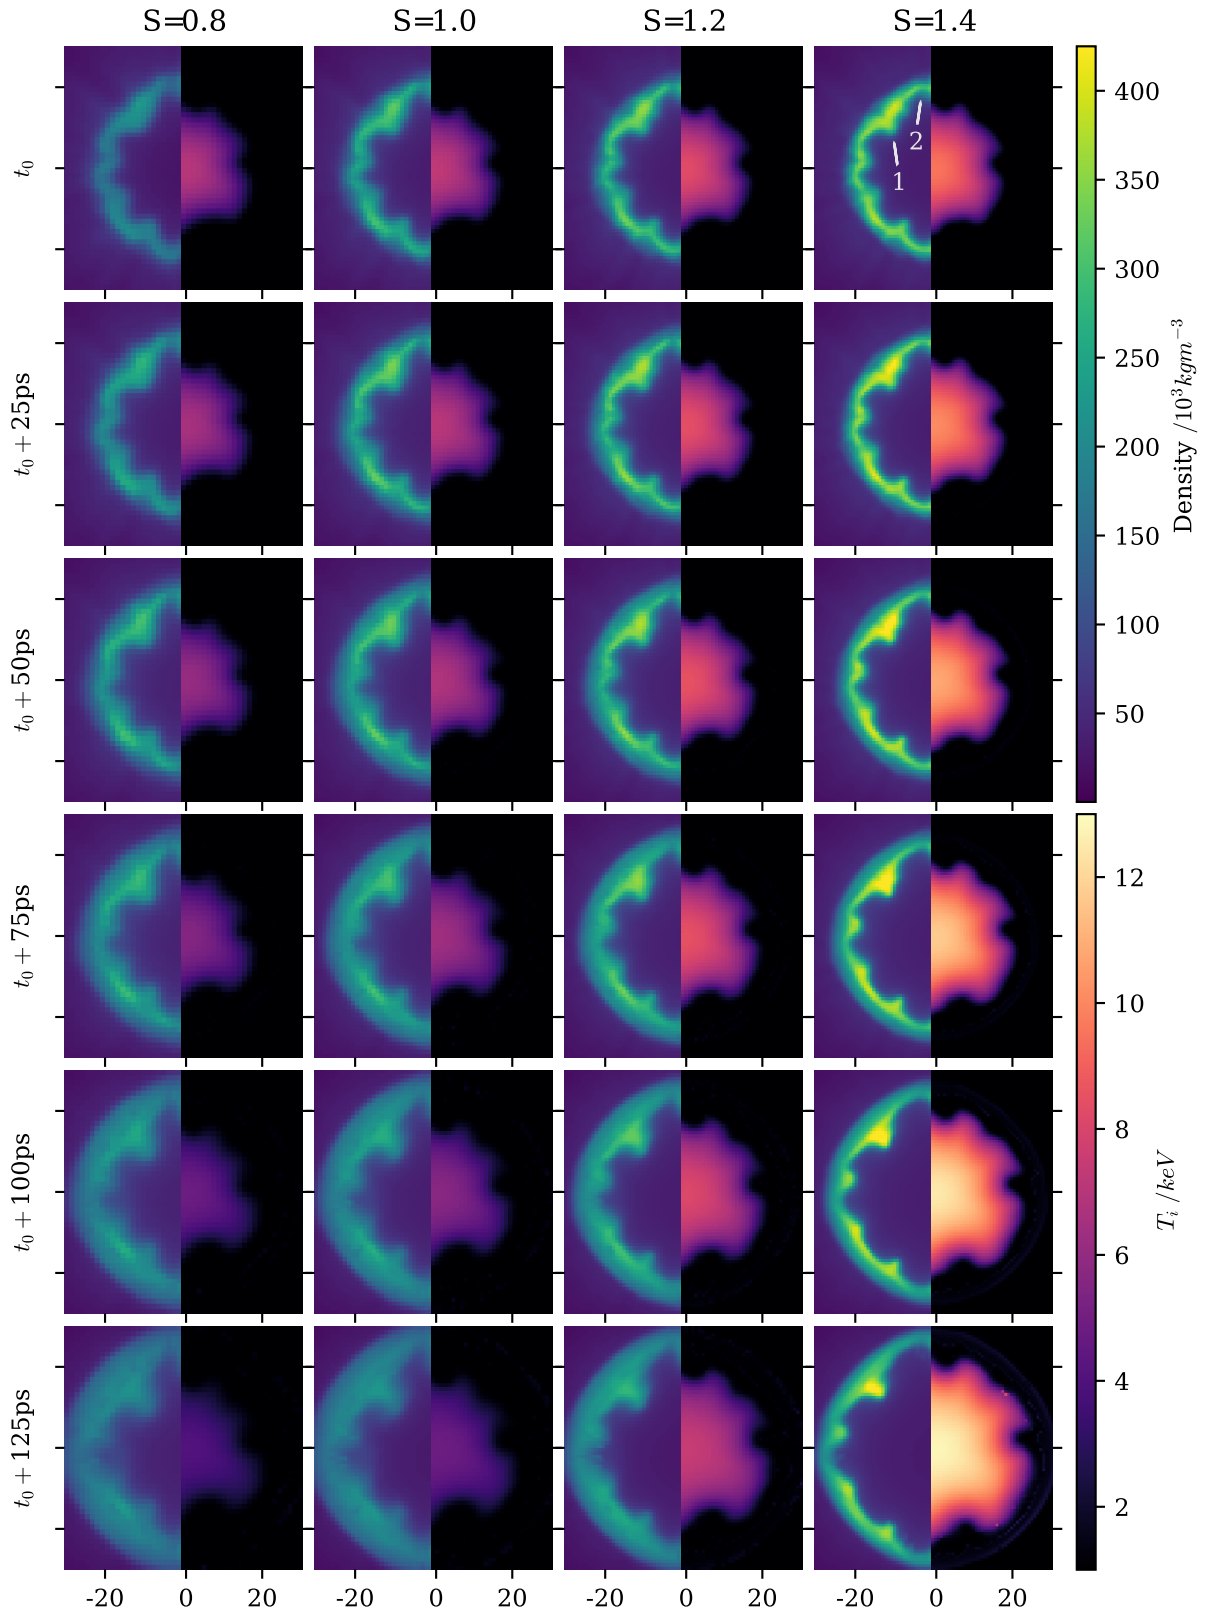
Figure 6.18: A grid showing the time-evolution (down the grid) of density (left half) and ion temperature (right half) slices for increasing scale factor, S across the grid for the multi-mode scenario. The physical scale is normalised, x/S (μm), in order to better compare the features across scale factors. Times are shown relative to the time of peak compression, t_0. Annotations '1' and '2' indicate features exhibiting perturbation ablative stabilisation and bubble expansion.
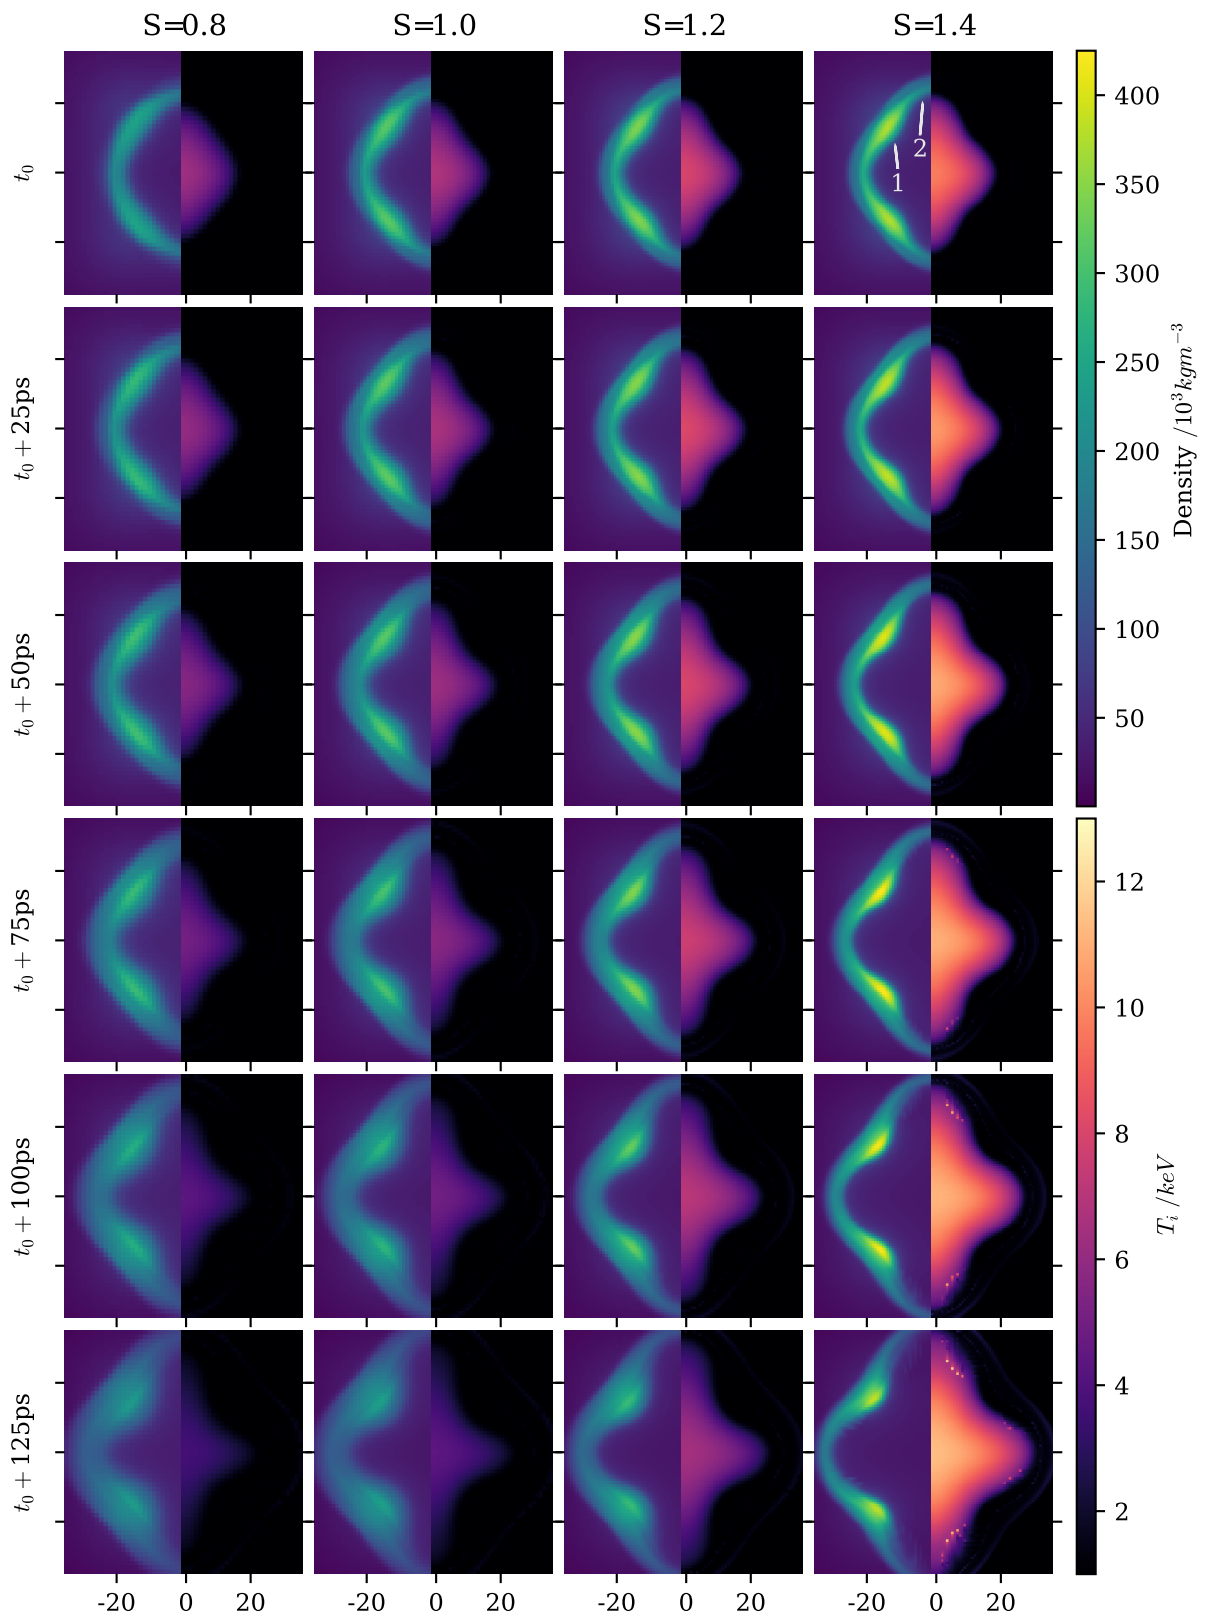
Figure 6.19: A time-evolution grid plot of density and ion temperature against scale factor as in figure 6.18, but for the radiation asymmetry scenario.
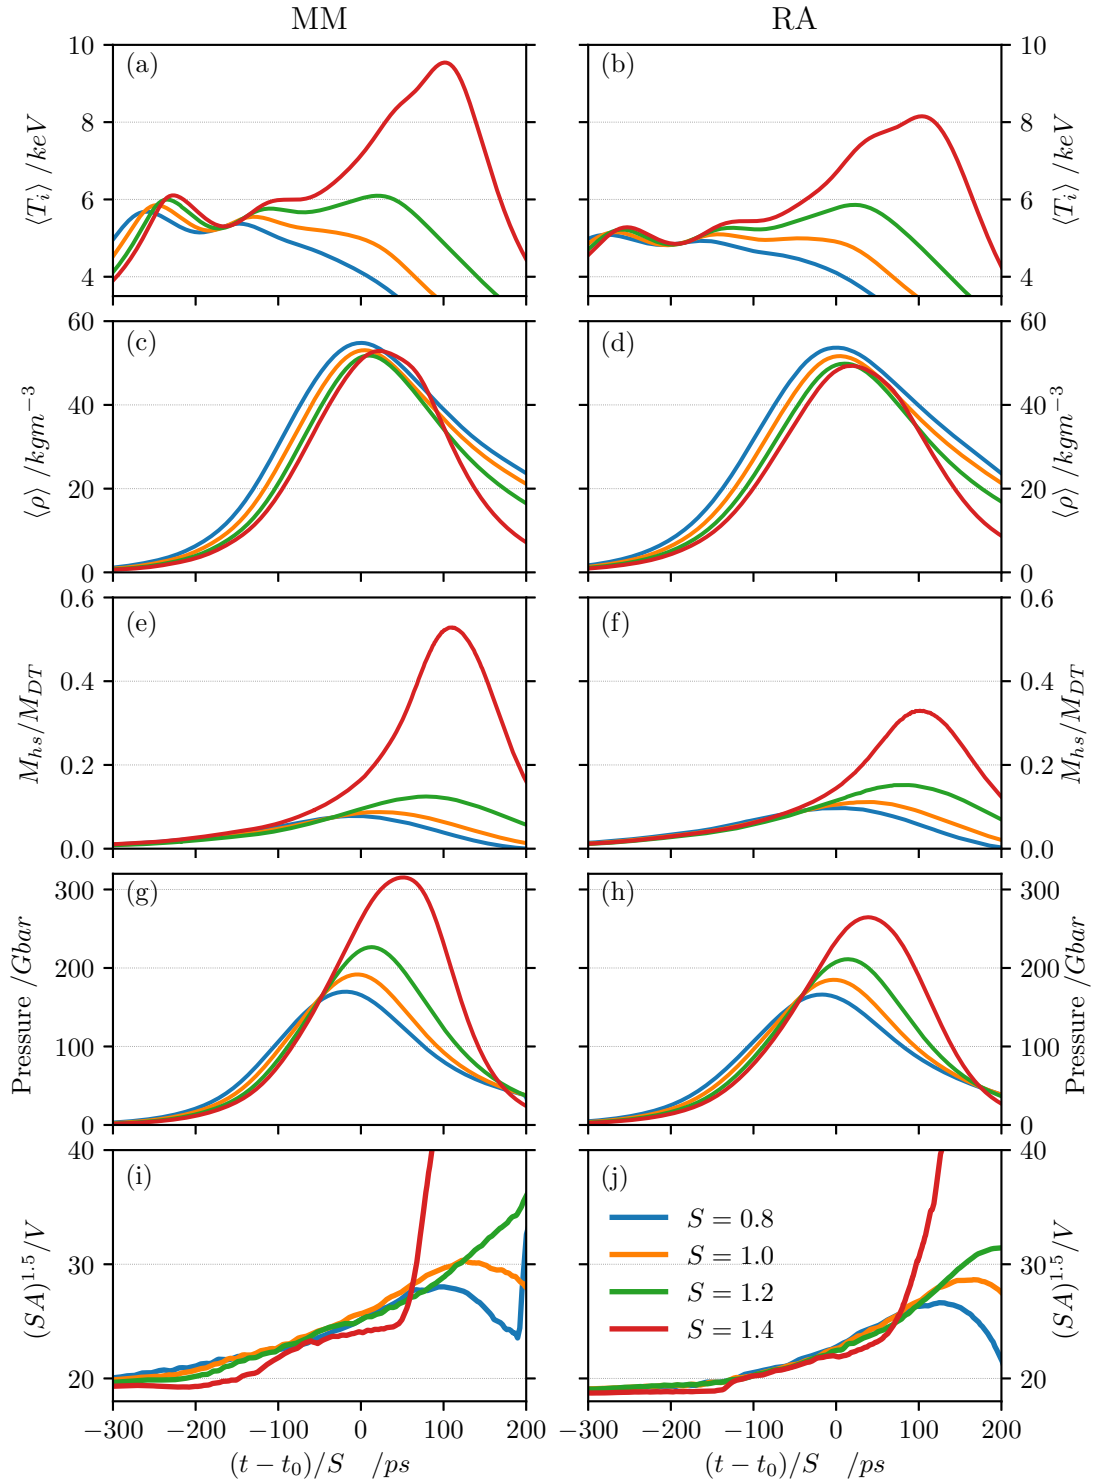
Figure 6.20: (a,b) Burn-averaged ion temperatures, $\langle T_i \rangle$; (c,d) burn-averaged density, $\langle \rho \rangle$; (e,f) mass relative to total DT fuel mass, M_{hs}/M_{DT}; (g,h) pressure, P; and (i,j) dimensionless surface area to volume ratio, $(SA)^{1.5}/V$ of the hotspot for MM (left) and RA (right) respectively, across all four scales: $S = 0.8$ (blue), $S = 1.0$ (yellow), $S = 1.2$ (red) and $S = 1.4$ (green). Times are normalised by the time of peak compression, t_0 and adjusted for scale factor S.

The hotspot temperatures seem to roughly agree with $T \sim S^{0.2}$ [128, 129] until $t \sim t_0 - 100$ps. Beyond this point, $\langle T_i \rangle$ diverges from this scaling for the higher scales as the alpha-heating starts to become significant, with the largest divergence for S=1.4. S=1.0 still seems to weakly agree with $T \sim S^{0.2}$ due to the relatively weak alpha-heating levels. The hotspot pressure does not remain constant in this scaling, even for the weak alpha-heating levels seen at S=0.8 & 1.0.

Here we note that the profiles in figures 6.20a and 6.20b remain slightly shifted as the times are scaled to the time of peak compression, rather than the time of the first shock-flash (which coincides with the first peak in the figures). The scaled time between the time of peak compression and the time of the first shock-flash reduces with scale factor as the alpha-heating increases the hotspot pressure earlier, resulting in a faster deceleration of the shell for higher scale factors. Although the pressures shown in figures 6.20g and 6.20h appear to indicate lower pressures at higher scale at this time, this is again due to the timings being relative to the time of peak compression.

The pressure serves as a good indication of the hotspot performance, with the strongest performers producing the highest peak pressure. The increase in pressure between S=1.2 and S=1.4 for MM is significantly larger than for RA, matching the faster yield scaling for MM. The peak pressure remains relatively similar between MM and RA at S=0.8 and S=1.0, with the corresponding yields (for each scale factor across the scenarios) also being similar. The non-hydrodynamic scaling of the pressure results in the same increased back-compression of the shell as seen in section 4.3. This can be seen most clearly comparing across the scale factors in figures 6.18 and 6.19, with the back-compression increasing the shell density and reducing the shell thickness.

The enhanced heat flow from the increased alpha-heating and thermal conduction at larger scale factors produces stronger shell ablation, evidenced by the increased mass accrual of the hotspot. This ablation pressure further contributes to the thinner and denser shell. The higher peak density is also a factor in sharpening the density gradient between the hotspot and the shell, and similarly the increased temperature gradient at the edge is influenced by the significant boost in the central temperature. The hotspot-shell boundary region is also thinner, with the hotspot and the shell each more strongly defined in this regime.

As initially seen in figure 6.16a, figure 6.20i and 6.20j show the ratio $(SA)^{1.5}/V$ is higher for MM than for RA across all scales.

The enhanced ablation results in stronger perturbation stabilisation and reduced growth due to fire-polishing, which can be readily seen at the regions indicated by '1' on figures 6.18 and 6.19. Visually, the perturbation penetrates less deeply and recedes faster. More quantitatively, this can be seen in the plots of $(SA)^{1.5}/V$ in figures 6.20i and 6.20j for S=1.4, and is clearest for MM due to the stronger burn. The ratio for S=1.4 MM is noticeably lower than for the lower scale factors, and in particular increases much slower between $(t - t_0)/S \sim -50$ps and $(t - t_0)/S \sim 50$ps. The reduced perturbation growth slows the increase in $(SA)^{1.5}/V$ and therefore results in a more spherical hotspot than for the lower scale factors without the fire-polishing.
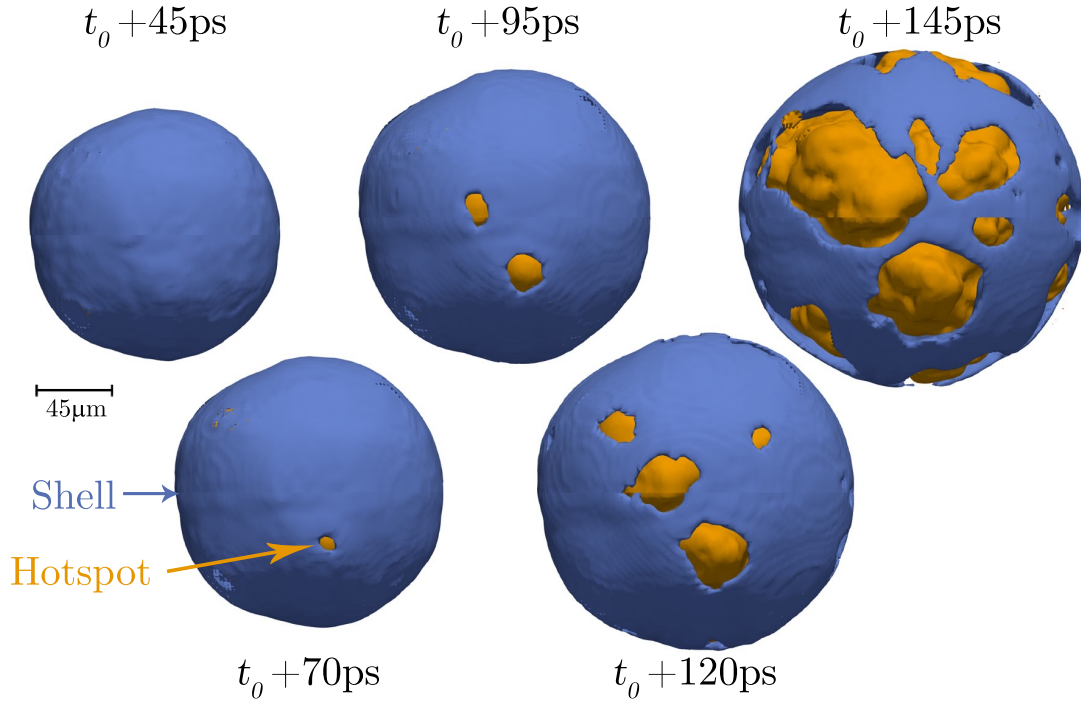
Figure 6.21: Time evolution of the MM scenario showing multiple "aneurysms" of the hotspot (shown here by the $T_i = 3\text{keV}$ contour in orange) bursting out through the shell (shown by the $\rho = 210,000\text{kgm}^{-3}$ contour in blue).

The increased hotspot pressures at larger scale factors also result in correspondingly faster hotspot re-expansion. Re-expansion into regions of weaker confinement (i.e. low ρR) will naturally occur faster than for regions of stronger confinement (high ρR) due to their lower inertia. The increased hotspot pressures exacerbate this differential expansion rate between the weak and strong confinement regions, and result in 'aneurysms', i.e. loss of confinement through these weak regions [39]. This loss of confinement results in a significant deviation in shape from spherical symmetry, and can be noticed by the particularly sharp increase in $(SA)^{1.5}/V$ at $(t-t_0)/S \sim 50\text{ps}$ for $S{=}1.4$ MM and RA in figures 6.20i and 6.20j. This increase coincides with the peak and subsequent drop-off in pressure seen in figures 6.20g and 6.20h, indicating a loss of hotspot pressure. The difference in expansion rates can be seen at the region annotated '2' in figures 6.18 and 6.19, which expands significantly faster than the high ρR region indicated by '1'. Note that this expansion in figures 6.18 and 6.19 is faster even in real time, not just scaled time — implosion timescales should be lengthened by S, and so hydrodynamic scaling would expect the expansion to be slower in real time, not faster.

This aneurysm-like feature can be clearly seen in figure 6.21, which shows the time-evolution of the hotspot and shell in the MM $S{=}1.4$ scenario. The hotspot is indicated by the orange $T_i = 3\text{keV}$ contour, and the shell by the blue $\rho = 210,000\text{kgm}^{-3}$ contour. The first time-step shown is just before the sudden increase in $(SA)^{1.5}/V$ shown in figure 6.20i, with the next time-steps showing the hotspot expanding into the lower density regions of the shell, which give weaker confinement.
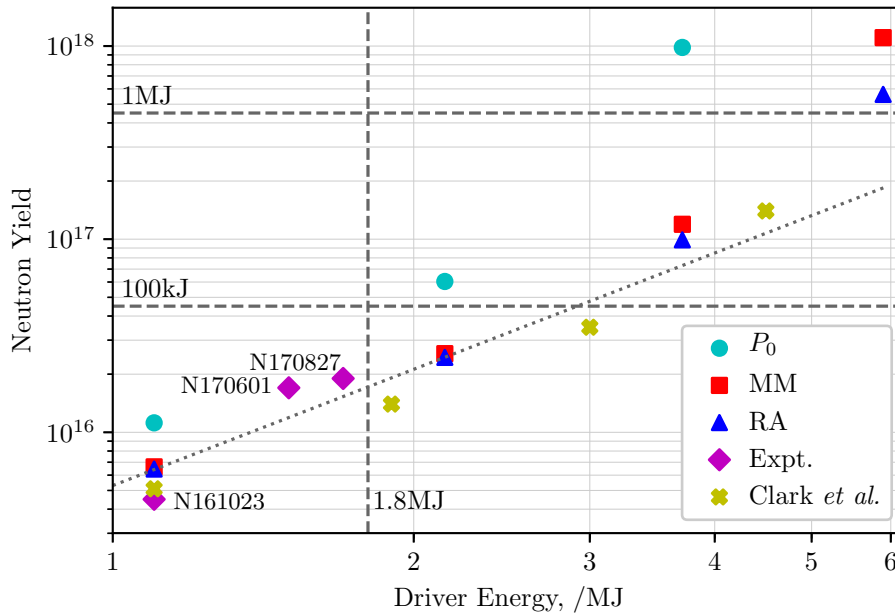
Figure 6.22: Yields as a function of driver energy shown on a log-log scale for the 3D simulations of the symmetric P_0 (cyan circle), the multi-mode (red square) and the radiation asymmetry (blue triangle) scenarios. The experimental yields of N161023 [3], N170601 and N170827 [9] are shown in magenta triangles, and Clark *et al*'s [3] simulation scaling of N161023 is shown in yellow crosses. The NIF laser energy of 1.8MJ is marked in dashed, as are energy yields of 100kJ and 1MJ.

The effect of the faster re-expansion and increased heat flow on the hotspot power balance can be seen from the power balances in figure 6.17d. The figure shows not just larger peaks in W_{net} with S due to the increasing alpha-heating levels, but also lower troughs, indicating stronger losses. The higher temperature raises the thermal conduction and radiative losses, and the increased density (or rather, increased duration of high density) from the shell mass ablation also affects radiative losses. Stronger expansion losses follow from the raised hotspot pressure, which induces faster re-expansion. We note that the aneurysm-like behaviour only occurs in the propagating burn regime. Lower alpha-heating regimes are either unable to generate the pressures required for this to occur, or do not possess the average shell areal density required to contain the significant hotspot pressure, such that there is little difference in re-expansion rate between the low ρR region and an average region of the shell.

Although we are scaling isolated perturbation scenarios rather than a comprehensive set of perturbations, it is still of interest to compare and contrast the results with other experimental and computational scaling studies. Figure 6.22 shows the yields of our 3D simulations, replotted as a function of driver energy, and compared to the yields of Clark *et al.*'s simulation study [3] and the experimental yields of N161023 [3], N170601 and N170827 [9]. Clark *et al.*'s scaling study explored the yield scaling of N161023 using 2D simulations with an extensive set of perturbations sources including

the tent, fill-tube, surface roughness and radiation drive asymmetries. These simulations achieve reasonable agreement with the experimental N161023 yield, and predict yields of $\sim 1.3 \times 10^{16}$ at full NIF laser energy of 1.8MJ.

We note that both of the MM and RA scenarios at S=0.8 overpredict the yield at $\sim 6.5 \times 10^{15}$ compared to the experimental yield of $\sim 4.5 \times 10^{15}$ [3]. The recent highly performing experiments produced yields of 1.7×10^{16} and 1.9×10^{16} [9], scaling up the capsule inner radius by 8% to $910\mu m$ from $844\mu m$ and the hohlraum to 6.2mm from 5.75mm. These experiments are not a direct hydrodynamic scaling, since these experiments were shot with laser energies of 1.5MJ and 1.7MJ respectively, rather than the 1.4MJ given by the direct S^3 scaling. The experimental yields improve due not only to the increased scale factor, but an extra increase in laser energy, reduced coast time [9] and a reduction in the fill-tube diameter to $5\mu m$ [57]. These indicate the kinds of experimental improvements that can greatly enhance any hydrodynamic scaling applied.

Meanwhile, interpolating between our data gives yields of $\sim 1.5 \times 10^{16}$ at S=0.9, and $\sim 2 \times 10^{16}$ at laser energies of 1.8MJ ($S \sim 0.95$). This is slightly greater than given by Clark *et al.*, but similar to the highest performing experiments.

The yield scaling of our scenarios agree reasonably well with Clark *et al.*'s, although our simulations begin from a higher starting yield. Our isolated perturbation scenarios would need to be scaled up to laser energies of ∼5-6MJ in order to achieve an energy yield of 1MJ. However, as seen from the experiments, and explored by Clark *et al.*, this can be significantly improved by reducing the impact of perturbations. We note that the perturbations in our scenarios increase in wavelength in tandem with the scenario scaling, but that this may not be the case for experiments. For example, the size of the fill-tube is likely to remain fixed, and the scale length of surface roughness perturbations will also remain fixed rather than scale up with size. The decreasing size of the perturbations relative to the capsule size is likely to reduce its impact and improve the overall scaling of experiments.

6.3 Summary

In this chapter, we have explored the hydrodynamic scaling of implosion designs based on both Highfoot N130927 and HDC N161023 by introducing yield degradation levels to match those seen on the respective experimental shots. In both scaling studies, two types of scenarios were explored; a short-wavelength multi-mode perturbation (MM) based on S. Taylor's [91] methodology for implementing RTIs in the velocity field using Layzer's approximate analytic treatment [186] (as in section 5.1.3), and a long-wavelength perturbation based on asymmetries in the radiation drive (RA).

In both the Highfoot scaling and the HDC scaling, the subtleties in the delivery of PdV work to the hotspot had a significant impact on the yield scaling. In the unperturbed, symmetric scenario, the PdV work is done in the shortest amount of time, as there is little variation in the shell velocity; all regions of the shell decelerate on the centre of the hotspot symmetrically, at roughly the same time. For perturbed scenarios

however, the PdV work is spread out over a longer period of time, as the variation in shell velocity across perturbed regions results in differing times of stagnation. The increased duration of the compression work being done on the hotspot allows more time for heat flow from the hotspot back into the shell, resulting in a less adiabatic compression. The resultant heating of the hotspot is weaker, producing lower temperatures at peak compression.

In the Highfoot study, although MM produces a lower yield than RA, it scales faster. Although the timescale of the PdV delivery, τ_m (equation 6.1), for MM is higher than for RA, it decreases faster and the total work done on the hotspot is also higher for MM as well. However, in general, both the MM and RA scenarios in the Highfoot study perform weakly (in yield). The significant yield degradation levels resulted in a low initial starting yield with a very small alpha-heating contribution. The hydrodynamic scaling did produce an increase in the alpha-heating, but the difference is not significant. At S=1.4 ($\hat{S} \sim 1.2$), the alpha-heating for both MM and RA deposited only $\sim 55\%$ and $\sim 80\%$ of the PdV work done on the hotspot, up from $\sim 30\%$ and $\sim 45\%$ respectively at $S = 1.0$ ($\hat{S} \sim 0.8$). As such, neither scenario in the Highfoot scaling produced an igniting hotspot.

For the HDC scenario, both perturbation scenarios exhibit igniting behaviour at the higher scale factors (S=1.2, 1.4), with curvature in the yield-scale graph in figure 6.13. In both scenarios, the yield degradation is much less than for the Highfoot scenario, and at the starting point of the scaling (S=0.8) the alpha-heating energy deposited in the hotspot is similar to the mechanical work deposited. The increase in scale factor raises the alpha-heating levels in the hotspot, because: *a)* the hotspot volume is larger and contains more fuel, producing more fusion reactions and therefore more alpha-particles; *b)* the larger hotspot radius increases the hotspot areal density, ρR, and therefore the hotspot also retains more of the alpha-heating; and *c)* the longer implosion timescale and confinement time allowing more time for fusion reactions to occur and for alpha-heating to bootstrap. The significant improvements to the symmetry of the HDC implosions result in a much better initial performance than for Highfoot, even at a reduced starting scale factor of S=0.8 ($\sim 6 \times 10^{15}$ for HDC compared to $\sim$ 3-4 $\times 10^{15}$ for Highfoot).

The differences between the two perturbation scenarios are more apparent with the strong alpha-heating levels, with the MM scenario scaling faster (in yield) than the RA scenario. This is due to more synchronous PdV compression for MM than RA producing higher temperatures, and therefore a stronger alpha-heating bootstrap at higher scale factors.

Between the Highfoot scaling study and the HDC scaling study, it is clear that increasing the scale factor of existing implosions will not overcome severe yield degradations to produce high yields (or rather, a very large increase in scale factor would be necessary). Improvements in perturbation levels are instead a more efficient way to achieve significant boosts in yield, particularly when combined with (smaller) increases in scale factor.

Our HDC scaling study also demonstrated behaviours of igniting and burning hotspots

unique to perturbed conditions. The burning hotspots at $S=1.4$ exhibited significant levels of fire-polishing due to the preferential heat flow into perturbations (as demonstrated in section 5.2) from both alpha-heating and thermal conduction. In addition, the combination of differences in areal density with high hotspot pressures in the propagating burn regime resulted in a significant difference in hotspot expansion rates between regions of weak confinement and regions of strong confinement. This produced "aneurysms", in which confinement is lost through regions of weak confinement and burn is truncated.

We also compared the results of our HDC scaling study to other examples of HDC implosion scaling, both experimental [9] and computational [3] (with a comprehensive set of perturbation sources). The isolated and idealised nature of our perturbation sources scaled up in size with the capsule scaling, whereas experimental perturbations do not necessarily follow this; for example, the size of the fill-tube perturbation will remain the same, even as the capsule size increases. Even so, our study exhibits similar curvature in the yield scaling to Clark *et al.*'s, albeit slightly higher in yield due to a less perturbed starting point. The experimental yields increase faster than both computational studies, due to an increase in implosion velocity (through larger increases in laser energy than dictated by hydrodynamic scaling, reductions in the fill-tube diameter, and reduced coast time) in addition to the increase in scale factor. Interpolating between our scenarios gives a requirement of 5-6MJ of laser energy to reach 1MJ of energy yield, while a similar extrapolation from Clark *et al.*'s result gives a similar laser energy. However, the improvements to experimental performance through reductions in the impact of perturbations have been seen to significantly affect the scaling, and therefore can drastically reduce this requirement.

7 Conclusions and Future Work

7.1 Conclusions and Further Research

A Monte-Carlo charged particle transport module has been developed for the radiation-hydrodynamics code Chimera, for the purpose of modelling alpha-heating in perturbed ICF hotspots. The model uses the Zimmerman formulation of the Maynard-Deutsch stopping model, and is implemented with a linked-list data structure that is ideal for this purpose. The implementation incorporates magnetised transport capabilities, which are not used in this work but open up avenues for future research.

1D Behaviour

1D simulations were used to gain an intuition for the behaviour of the hotspot and its power balance in chapter 4. Radiative transfer plays a non-negligible role in the hotspot power dynamics, with the hotspot reabsorbing a significant proportion of emitted radiation within the hotspot itself, in addition to absorption in the shell as well. This is particularly important in lower-yield regimes, in which the alpha-heating contribution is less significant in the overall hotspot power balance, and the hotspot is colder and therefore less transparent to lower energy photon emission.

Three regimes of alpha-heating were explored through hydrodynamic scaling with capsule size and driver energy; self-heating, robust ignition and propagating burn. Each regime has a number of defining characteristics, namely the mechanism of yield amplification, the strength of the alpha-heating contribution relative to the other contributions, and the time-variation in the burn-averaged hotspot temperature. The self-heating regime exhibits minimal yield amplification due to alpha-heating, which increases slower than the combined power losses from the hotspot and therefore results in a falling hotspot temperature at peak compression. Robust ignition yield amplification occurs due to significant increases in hotspot temperature, with the alpha-heating increasing at a similar rate to the power losses, resulting in the hotspot temperature remaining roughly constant at peak compression even as the capsule expands. Propagating burn amplifies the yield through not only increased temperatures but also through a significant increase in the total mass of fusing material ablated by the deflagration wave. Here, the alpha-heating increases faster than the hotspot power losses, such that the hotspot temperature continues to increase after peak compression.

These regimes correspond well to the ignition criteria described by Springer *et al.* [85] and Cheng *et al.* [87]. We find that the propagation of burn is dependent on the existence of significant levels of thermonuclear burn within the hotspot itself, as

considered in Cheng *et al.*'s model. Our results also corroborate the hotspot power balance considerations on which Springer *et al.*'s ignition criterion is based. In addition, the simulations indicate that yields of $\sim 10^{17}$ could be achieved within the robust ignition regime without necessarily entering the propagating burn regime, given the two stages of yield amplification from alpha-heating.

Alpha-heating significantly alters the hydrodynamics of igniting hotspots, with faster re-expansion, larger hotspots and more compressed shells due to the increased hotspot pressure, in addition to sharper temperature- and density-gradients. Such effects will be important to keep in mind in the design of capsules aiming to reach the propagating burn regime.

The impact of how the hotspot is defined on calculated hotspot parameters was explored. The subject is particularly important in low-yield regimes, where the distinction between the hotspot and the shell regions is poorly defined, with large temperature- and density-gradient scale lengths. We note that the experimentally observable 17% neutron contour can be significantly smaller than other commonly used definitions, such as ion temperature or density contours. These definitions do not adapt across alpha-heating regimes, while intuitive definitions such as the 'heat-source' of the capsule are dynamic in this regard. However, such definitions can be difficult to extract and it is unclear that any particular definition is more appropriate to use than any other in general. Rather, the manner in which the hotspot is defined has a significant impact on quantities such as the hotspot burn-averaged temperature, particularly in low-yield regimes, and thus it is important to know the definition used in order to make fair comparisons. This is particularly salient given the importance of these parameters in comparing experiments to simulations, and the self-heating regime in which the experiments are performing.

3D Idealised Perturbations

Following the insights gained in 1D, the impact of idealised perturbations on the hotspot was explored using 3D simulations in chapter 5. The spike perturbation caused a significant increase in heat flow out of the hotspot, through mechanisms of both thermal conduction and alpha-heating. The thermal conduction increases as the temperature gradient around the spike is larger, in addition to the increase in surface area due to the spike's distortion. The displacement of cold dense shell material towards the region of greatest fusion production results in a larger flux of alpha-particles into the spike. These alpha-particles are also more energetic at the point of encountering the shell material, and thus more alpha-heating is deposited into the spike. Furthermore, the perturbation results in less efficient conversion of shell kinetic energy to thermal energy, which has a significant knock-on impact on the strength of the alpha-heating feedback loop. The bubble perturbation similarly reduces the efficiency of PdV heating of the hotspot, but does not increase the heat flow out of the hotspot, and therefore has a much lower impact on the ignition of a capsule than an equivalent spike perturbation. However, the low areal-density of the bubble can compromise the confinement of the hotspot and

therefore truncate the propagation of burn.

We therefore expect the primary impact of spike perturbations to be the degradation of the formation and heating of the hotspot within the ignition process, whereas bubble perturbations instead degrade the hotspot confinement and interfere with the burn propagation. Further investigations could study this by exploring the impact of the spike and bubble perturbations at different scale factors, and therefore in different alpha-heating regimes. Since the primary effect of the bubble perturbation is to cause loss of confinement, it is expected that the impact of a bubble perturbation would be minimal for self-heating, minor for robust ignition and significant in the propagating burn case. Meanwhile, the impact of a spike perturbation would likely produce a notable reduction in performance for both the self-heating and the robust ignition regimes, as was seen in this work. In the propagating burn regime, the spike would reduce the strength of the bootstrap and make ignition more difficult. However, it would also be subject to fire-polishing and would not reduce the confinement as a bubble perturbation would, and therefore the overall effect of such a perturbation is difficult to estimate.

Such an investigation could be further extended by exploring the wavelength- and amplitude-dependence of the impact. We anticipate that a shorter wavelength would reduce the impact; a shorter wavelength spike is easier to stabilise via ablation, while a shorter wavelength bubble equates to a smaller region of low areal density, reducing the size of the region of weak confinement. Comparisons of the amplitude's effect on the performance impact could be made if the amplitudes were matched using a cross-compatible methodology.

These idealised perturbations could also be combined to explore how perturbations might interact with one another in a simplified scenario. For example, would a spike directly opposite a bubble induce faster loss of confinement? Or would such a spike reduce the performance of the capsule and thus reduce the hotspot pressure, such that the bubble expansion is reduced?

The difficulty in reaching the propagating burn regime through increasing the implosion velocity alone was highlighted; the increase in kinetic energy delivered to the hotspot increases the temperature and thus the hotspot is more likely to ignite robustly, but the lack of a corresponding increase in areal density (and therefore confinement) results in much faster re-expansion of the hotspot. The confinement must be adequately increased in conjunction with improving the heating of the hotspot in order for a burn wave to propagate.

3D Perturbed Scaling

Chapter 6 explored the hydrodynamic scaling of perturbed hotspots based on both High-foot N130927 and HDC N161023 implosion designs, using short-wavelength multi-mode and long-wavelength radiation asymmetry perturbations. Variations in the synchronisation of PdV delivery between the perturbations were found to affect the scaling, since an increase in the time duration over which PdV work is done also allows more time for heat

to flow from the hotspot into the shell. The short-wavelength multi-mode perturbation scenario scaled faster in yield than the long-wavelength radiation asymmetry in both implosion designs. This was in part due to a lack of mass redistribution in the perturbation scenario, resulting in smaller variations in shell momentum and therefore better synchronisation of PdV work. Accordingly, this suggests that improvements to long-wavelength asymmetries may be more efficient than improvements to short-wavelength asymmetries in the effort to increase yield through hydrodynamic scaling.

The igniting perturbed hotspots also demonstrated significant fire-polishing driven by alpha-heating and thermal conduction, in addition to "aneurysm"-like loss of confinement due to variations in shell areal density. The combination of high hotspot pressure with sufficiently large variations in confinement resulted in a significant enough difference in expansion rate between these regions to induce such an effect. Although such effects are unlikely to be limiting the performance of current experiments, they are likely to become increasingly important as performance improves.

The yield degradation level had a significant impact on the efficacy of using hydrodynamic scaling to increase the yield. The relatively weak alpha-heating levels in N130927 due to significant yield degradation resulted in only a linear yield scaling with minor increases in the alpha-heating level. N161023 had a significantly lower yield degradation level, and thus hydrodynamic scaling resulted in both perturbation scenarios exhibiting non-linear yield scaling and demonstrating characteristics of the robust ignition and propagating burn regimes by $S = 1.4$.

Our N161023 perturbation scenarios also exhibited similar yield scaling curvature to that shown in other work by Clark *et al.* [3]. This is despite the fact that the size of our isolated perturbation scenarios scaled directly with the capsule scaling, whereas the realistic perturbations used by Clark *et al.* — such as the tent and the fill-tube — do not, and therefore the impacts of which should decrease. Predictions from both studies produce requirements of 5-6MJ of laser energy to reach 1MJ of yield.

However, it is clear from the comparison between the N130927 scaling and the N161023 scaling that increasing the scale factor alone is not the most efficient path towards achieving megajoule yields. Rather, mitigating the yield degradation mechanisms is a much more efficient way to achieve significant yield gains, particularly when combined with hydrodynamic scaling. This is also made clear from the improvements in experimental performance, which have a better yield scaling than suggested by experiments precisely because of this combination of reducing yield degradation mechanisms (such as from reducing the size of the fill-tube [57]) with increasing the scale factor and laser energy. Such results suggest a positive outlook that the above requirement can be significantly reduced.

Indeed, some perturbation effects may decrease with scale anyway — for example, the size of the fill-tube will not increase with scale factor, and therefore the relative impact ought to decrease. Although our perturbation scenarios are isolated, interactions are likely to occur between the multiple sources of perturbations, and it is as yet unclear how these interactions would scale. With work currently ongoing within the research group to

accurately model the impacts of the tent-scar, surface roughness, fill-tube and radiation asymmetries for N170601, these perturbations could allow for further development of this scaling study, both as individual perturbations and as a combination.

An interesting scenario to explore would be of two types of perturbation set up to mitigate the effects of one another. Since individually the perturbations would likely scale differently, would the scaled versions still have the same mitigating interaction? How would each perturbation have to be modified in order to maintain said interaction? Such questions are particularly interesting in light of the fire-polishing and aneurysm-like behaviour exhibited by highly-performing perturbed capsules. Increasing the types of available perturbations would allow us to explore numerous scenarios of interactions between perturbations, and how the reductions in certain perturbations could affect the yield in concert with hydrodynamic scaling.

7.2 Code Improvements

There are a number of ways in which the charged particle transport module could be upgraded. The parallelisation scheme could be improved to better balance the computational load across processors. The current scheme follows the domain decomposition of Chimera. However, the use to simulate the ICF hotspot means that only the processors containing a portion of the hotspot are used to model the alpha-particles, while the processors containing the outer regions of the capsule do not contribute to the alpha-particle modelling. An alternative scheme could be to pass the entire domain in which alpha-heating is significant to all processors, and then have each processor simulate alpha-particles across the entire domain. This domain could either be predetermined at initialisation or dynamically adjusted – the former would be much less flexible, but also have a lower computational overhead. A similar scheme has been implemented in work using an adapted version of the module to model planetary radiation belts [192].

It is currently not necessary to include the impact of thermalised alpha-particles (helium ash) on the hydrodynamics of the plasma, as the number density of alpha-particles is several orders of magnitude below that of the DT ions due to the low burn-up fraction of the capsule. However, in designs with larger burn-up fractions, the impact of the helium ash on the hydrodynamics will become much more significant and will therefore need to be implemented. Fortunately, recent work in the group has generalised the number of materials in Chimera, and thus allows for a relatively simple implementation of helium ash.

The alpha-particle model is currently implemented in the Cartesian and spherical geometries of Chimera. There are future plans to port the model into cylindrical geometry, which, in combination with the magnetised particle transport capabilities and ongoing work to port the existing extended MHD capabilities [131] into cylindrical geometry, would allow the exploration of ignition and burn in MagLIF style implosions.

7.3 Future Work

The inclusion of magnetised particle transport opens the way to exploring the growing area of research in magnetised ICF [27, 193], tying in with existing research avenues in the research group into magnetic fields in ICF implosions [131, 193]. Indeed, Walsh *et al.* [193] have already begun to explore the interactions of magnetisation and perturbations with magnetised alpha-particles using this module. Of particular interest from the viewpoint of the current work is the impact of magnetisation on the processes of ignition and burn. Increasing magnetisation will progressively restrict the electron thermal conduction and the alpha-heating to smaller regions of the capsule, reducing heat loss from the hotspot and enabling ignition. However, the propagation of burn is dependent on the electron thermal conduction and alpha-particle transport to transport heat flow away from the hotspot into the shell. By inhibiting these heat flow mechanisms, the ignition threshold is lowered but burn propagation is restricted, and it is unclear how this might evolve in a 3D perturbed system with complex magnetic field topologies. It is possible that the magnetic containment of the heat within the hotspot could generate a detonation wave due to the growth in pressure in the hotspot. The anisotropy in the magnetic field could allow for preferential burn propagation along the polar directions, in which case the capsule could be designed with more material in these directions to feed the burn.

The charged particle transport module developed in the course of this work, although applied to modelling alpha-particles, is generalisable to other species. This could potentially be used in the investigation of problems such as electron pre-heat, which is thought to affect the capsule compressibility and increase the adiabat of implosions [72, 194]. The capabilities for modelling the energy deposition of fast ions also potentially allow for the simulation of fast ignition implosions. The underlying computational methods provide the backbone for the potential development of other Monte-Carlo simulation tools, for purposes such as radiation transport or neutron transport. However, this is unlikely to be necessary in the near future due to existing tools for modelling these phenomena, and the computational expense of the Monte-Carlo method with the short timesteps required in both cases.

The module could also be applied to improving synthetic diagnostic tools [195] such as proton radiography [196]. Previous work has included charged-particle tracking with magnetisation and Spitzer slowing in a post-process scenario, but more advanced stopping models and larger numbers of particles would improve comparisons with experimental results. The spectroscopy of $14.7 MeV$ protons produced in $D + {}^3He \rightarrow T + p$ reactions has been used to reveal areal density asymmetries in OMEGA experiments [197], while similar work has also been done on OMEGA using spectroscopy of knock-on deuterons [198]. Here, neutrons from DT fusion reactions scatter elastically off deuterons in the dense shell, producing knock-on deuterons with energies up to 12.5MeV. The slowing of these deuterons affects the spectrum, and can reveal information about the plasma conditions and areal density. Another interesting diagnostic could be to model the fast

tritons produced in DD implosions ($D + D \rightarrow T + p$), which can react with D to produce secondary DT neutrons. In a magnetised scenario (such as in a MagLIF implosion), the resultant neutron spectra is anisotropic as a result of the magnetisation of the tritons [199].

Having demonstrated the different characteristics and hydrodynamic behaviour of different burn regimes, is it possible that different burn regimes might exhibit different diagnostic signatures? Although the yield is likely a good indicator of this, such signatures might prove useful if they can distinguish between the different yield amplification mechanisms, allowing a deeper insight into the behaviour of an experiment. Given that the propagating burn regime affects the shell properties, causing it to be denser, thinner and expanding faster than in either of the robust ignition or self-heating regimes, information about the shell conditions could elaborate on the regime of an experiment. One method of doing this could be to use reaction-in-flight (RIF) neutron spectra; here, deuterons and tritons knocked-on by the process mentioned above undergo a DT reaction in-flight, producing a significantly higher energy neutron which can be detected separately. The spectra of these neutrons will be dependent on the deuteron energy spectrum, and therefore the stopping power at the plasma conditions of the shell. Hayes *et al.* [108] used RIF neutron spectra and plasma conditions from simulations to test stopping powers in degenerate plasmas. Here, the RIF spectra could instead be used to infer the plasma conditions, as was done by Cerjan *et al.* [200]. The stopping model dependence could be constrained through the use and comparison of multiple stopping models.

Appendices

A Radiation Transport

Radiation transport describes how photons propagate through and interact with materials.

A.1 Radiative Transfer Equation

Considering first the radiation without any sources or sinks, then the intensity of the radiation as it propagates through space and time will not change. Thus for radiation of intensity I_ν initially at position $\mathbf{r}$ traveling in direction $\mathbf{\Omega}$ at time t, then after a time interval τ the radiation will have propagated a distance $c\tau$ in direction $\mathbf{\Omega}$. The intensity remains unchanged, therefore $I_\nu(\mathbf{r} + \mathbf{\Omega}c\tau, \mathbf{\Omega}, t + \tau) = I_\nu(\mathbf{r}, \mathbf{\Omega}, \tau)$. Taylor expanding about $\mathbf{r}$ and t, and discarding second-order or higher terms, we then get:

$$\frac{1}{c}\frac{\partial I_\nu}{\partial t} + \mathbf{\Omega} \cdot \nabla I_\nu = 0 \tag{A.1}$$

However, as radiation propagates through a medium, the radiation intensity will be affected by absorption and emission within the medium. Radiation propagating through a distance $dr = cdt$ in a medium of absorptivity k_ν, change in intensity is given by:

$$dI_\nu = -k_\nu I_\nu dr$$

The absorptivity k_ν can be thought of as the probability per unit length that a photon of frequency ν will interact with the medium. The optical depth τ_ν is defined as $\int k_\nu dr$, and is the exponent in Beer's law:

$$I_\nu = I_\nu^0 e^{-\tau_\nu} \tag{A.2}$$

The amount of energy added to the radiation as it propagates through this distance dr of the medium is given by:

$$dI_\nu = +j_\nu dr$$

where j_ν is the emissivity. The emissivity describes the energy added to the radiation as it propagates through the medium, and the absorptivity the manner in which radiation is absorbed during said propagation. The absorptivity and spontaneous emissivity of a material are determined by the atomic characteristics of the medium, and thus vary depending on, for example, the ionisation state, density, and temperature of the medium. The calculation of absorptivities and emissivities falls within atomic physics, a separate topic unto itself, and is beyond the scope of this work. We note that the stimulated

emission should be mentioned separately, since the probability of the material emitting a photon due to stimulated emission is proportional to the intensity of radiation of the same frequency and direction as the emitted photon, i.e. the stimulated emissivity $j_{\nu,stim} \propto I_\nu$. Scattering processes also contribute to both the absorptivity and emissivity, since the photons are scattered to different angles and (sometimes) frequencies.

In general then, accounting for the absorption and emission of radiation, equation A.1 becomes the general radiative transfer equation:

$$\frac{1}{c}\frac{\partial I_\nu}{\partial t} + \mathbf{\Omega} \cdot \nabla I_\nu = j_\nu - k_\nu I_\nu \tag{A.3}$$

where j_ν includes both spontaneous, stimulated and scattering contributions to the emissivity, and k_ν includes both absorption and scattering contributions. It is important to note that k_ν and j_ν may be anisotropic, for example due to a preferred orientation of particles within the medium. Even if the absorption and emission were isotropic in the rest frame of the material, the resultant lab frame emissions and absorptions will not be due to the Doppler and aberration effects.

A.2 Radiation Moments

The angular moments of I_ν are the spectral radiation density;

$$U_\nu = c^{-1} \int_{4\pi} I_\nu d\Omega \tag{A.4}$$

the vector flux;

$$\mathbf{F}_\nu = \int_{4\pi} \mathbf{\Omega} I_\nu d\Omega \tag{A.5}$$

and the radiation pressure tensor:

$$\mathsf{P}_\nu = c^{-1} \int_{4\pi} \mathbf{\Omega}\mathbf{\Omega} I_\nu d\Omega \tag{A.6}$$

Taking the first and second angle moments of equation A.3 (i.e. multiplying by $\mathbf{\Omega}^0, \mathbf{\Omega}^1$ and integrating over $d\Omega$) we can obtain:

$$\frac{\partial U_\nu}{\partial t} + \nabla \cdot \mathbf{F}_\nu = 4\pi j_\nu - k_\nu c U_\nu \tag{A.7}$$

$$\frac{1}{c}\frac{\partial \mathbf{F}_\nu}{\partial t} + c\nabla \cdot \mathsf{P}_\nu = -k_\nu \mathbf{F}_\nu \tag{A.8}$$

since $\mathbf{\Omega}$ is constant. The first of these, equation A.7 represents the conservation of energy (or relative mass) for radiation of a given frequency, analogous to the first moment equation for fluids representing the conservation of mass. Similarly, the equation A.8 gives the conservation of radiation momentum, again for a given frequency.

Note that analogous to taking moments of the Boltzmann equation, we can keep taking moments but each successive moment introduces the next higher moment of

the intensity (or distribution function), and therefore there are always one too few equations to solve for the number of unknowns. In order to close the equations, another external relation between the highest moment and the lower moments must be found or approximated.

A.3 Approximations

A.3.1 Diffusion Approximation

The diffusion approximation is valid when the mean free path of the photon is small relative to other length scales, in other words in the limit of high opacity/absorptivity. Rearranging equation A.3 and expanding in $\frac{1}{k_\nu}$ gives:

$$I_\nu = \frac{j_\nu}{k_\nu} - \frac{1}{k_\nu}\left(\frac{1}{c}\frac{\partial I_\nu}{\partial t} + \mathbf{\Omega}\cdot\nabla I_\nu\right) \tag{A.9}$$

Assuming k_ν is large, the second term is a small correction to an isotropic radiation intensity, such that $I_\nu^0 = \frac{j_\nu}{k_\nu}$. Substituting this into the RHS of the above equation, we then get:

$$I_\nu^1 = \frac{j_\nu}{k_\nu} - \frac{1}{k_\nu}\left[\frac{1}{c}\frac{\partial \frac{j_\nu}{k_\nu}}{\partial t} + \mathbf{\Omega}\cdot\nabla\frac{j_\nu}{k_\nu}\right] \tag{A.10}$$

This then gives the following for the energy density and flux:

$$U_\nu = \frac{4\pi}{c}\frac{j_\nu}{k_\nu} - \frac{4\pi}{k_\nu c^2}\frac{\partial}{\partial t}\left(\frac{j_\nu}{k_\nu}\right) \tag{A.11}$$

$$\mathbf{F}_\nu = -\frac{4\pi}{3k_\nu}\nabla\frac{j_\nu}{k_\nu} \tag{A.12}$$

Equation A.12 is simply the diffusion equation for photons with diffusion coefficient $D = \frac{4\pi}{3k_\nu}$.

A.3.2 Eddington and $P_{1/3}$ Approximation

This then gives the pressure moment:

$$\mathsf{P}_\nu = \frac{1}{3}\mathsf{I}\left[\frac{4\pi}{c}\frac{j_\nu}{k_\nu} - \frac{4\pi}{k_\nu c^2}\frac{\partial}{\partial t}\left(\frac{j_\nu}{k_\nu}\right)\right] \tag{A.13}$$

The pressure moment from the diffusion approximation is isotropic, even though the intensity itself is not. This leads to the Eddington approximation:

$$\mathsf{P}_\nu = \frac{1}{3}U_\nu\mathsf{I} \tag{A.14}$$

Note that this is separate and distinct from diffusion. The Eddington approximation is weaker (i.e. less approximate) than the diffusion approximation, and therefore can be

used to close the equations A.7 and A.8 to get:

$$\frac{\partial U_\nu}{\partial t} + \nabla \cdot \mathbf{F}_\nu = 4\pi j_\nu - k_\nu c U_\nu \tag{A.15}$$

$$\frac{1}{c}\frac{\partial \mathbf{F}_\nu}{\partial t} + \frac{c}{3}\nabla U_\nu = -k_\nu \mathbf{F}_\nu \tag{A.16}$$

instead of equation A.8. Together the equations A.8 and A.16 form what is known as the P1 solution. However, if we drop the terms on the RHS of equation A.7, and combine the time-derivative with the divergence of equation A.16, then we get the following wave equation for the energy density:

$$\frac{\partial^2 U_\nu}{\partial t^2} - \frac{c^2}{3}\nabla^2 U_\nu = 0 \tag{A.17}$$

with wave-speed $c/\sqrt{3}$. Clearly, this is incorrect — and can be avoided by introducing a factor of 3 into the denominator of the derivative:

$$\frac{1}{3c}\frac{\partial \mathbf{F}_\nu}{\partial t} + \frac{c}{3}\nabla U_\nu = -k_\nu \mathbf{F}_\nu \tag{A.18}$$

which is known as the $P_{1/3}$ approximation [140]. This gives the same accuracy as the P_1 approximation in the optically thick limit, but also gives the correct solution in the limit of free-streaming radiation.

A.3.3 Coupling to Fluid equations

The radiation-matter interactions are all included in the source and sink terms of the radiative transfer equation, i.e. $j_\nu - k_\nu I_\nu$, and hence integrating the zeroth and first moments of this over frequency and angles will give the rates for energy and momentum exchange between a fluid and radiation. This is equivalent to integrating Eqns. A.7 and A.8/c over frequency, and inserting their negatives into the RHS of the material equations. Ignoring other contributions to the energy and momentum, this then gives:

$$\frac{\partial}{\partial t}\left(\rho e + U + \frac{1}{2}\rho u^2\right) + \nabla \cdot \left(\rho \mathbf{u}h + \frac{1}{2}\rho \mathbf{u}u^2 + \mathbf{F}\right) = 0 \tag{A.19}$$

$$\frac{\partial}{\partial t}\left(\rho \mathbf{u} + \frac{\mathbf{F}}{c^2}\right) + \nabla \cdot (\rho \mathbf{u}\mathbf{u} + \mathsf{P}) + \nabla p = 0 \tag{A.20}$$

for fluid density ρ, fluid velocity $\boldsymbol{u}$, fluid internal energy e, fluid pressure p, radiation energy density U, radiation flux density $\mathbf{F}$ and radiation pressure P.

B Figure Permissions

Figure Number	Reference	Copyright
Figure 1.1	[1]	AIP Publishing
Figure 1.3	[2]	AIP Publishing
Figure 1.3	[3]	IOP Publishing
Figure 1.4	[4]	IOP Publishing
Figure 1.5	[5]	AIP Publishing
Figure 2.1	[6]	AIP Publishing
Figure 3.8	[7]	Oxford Publishing Limited
Figure 6.2	[8]	AIP Publishing

Table B.1: List of permissions for figures reproduced in this thesis.

AIP PUBLISHING LICENSE
TERMS AND CONDITIONS

Mar 25, 2019

This Agreement between Mr. Jon Tong ("You") and AIP Publishing ("AIP Publishing") consists of your license details and the terms and conditions provided by AIP Publishing and Copyright Clearance Center.

License Number	4555960615336
License date	Mar 25, 2019
Licensed Content Publisher	AIP Publishing
Licensed Content Publication	Physics of Plasmas
Licensed Content Title	Progress towards ignition on the National Ignition Facility
Licensed Content Author	M. J. Edwards, P. K. Patel, J. D. Lindl, et al
Licensed Content Date	Jul 1, 2013
Licensed Content Volume	20
Licensed Content Issue	7
Type of Use	Thesis/Dissertation
Requestor type	Student
Format	Print and electronic
Portion	Figure/Table
Number of figures/tables	1
Title of your thesis / dissertation	Ignition and Burn in Perturbed Inertial Confinement Fusion Hotspots
Expected completion date	Apr 2019
Estimated size (number of pages)	200
Requestor Location	Mr. Jon Tong Blackett Laboratory Prince Consort Road London, London SW7 2BW United Kingdom Attn: Mr. Jon Tong
Total	0.00 GBP

Terms and Conditions

AIP Publishing -- Terms and Conditions: Permissions Uses

AIP Publishing hereby grants to you the non-exclusive right and license to use and/or distribute the Material according to the use specified in your order, on a one-time basis, for the specified term, with a maximum distribution equal to the number that you have ordered. Any links or other content accompanying the Material are not the subject of this license.

1. You agree to include the following copyright and permission notice with the reproduction of the Material:"Reprinted from [FULL CITATION], with the permission of AIP Publishing." For an article, the credit line and permission notice must be printed on the first page of the article or book chapter. For photographs, covers, or tables, the notice may appear with the Material, in a footnote, or in the reference list.

2. If you have licensed reuse of a figure, photograph, cover, or table, it is your responsibility to ensure that the material is original to AIP Publishing and does not contain the copyright of another entity, and that the copyright notice of the figure, photograph, cover, or table does not indicate that it was reprinted by AIP Publishing, with permission, from another source. Under no circumstances does AIP Publishing purport or intend to grant permission

to reuse material to which it does not hold appropriate rights.
You may not alter or modify the Material in any manner. You may translate the Material into another language only if you have licensed translation rights. You may not use the Material for promotional purposes.

3. The foregoing license shall not take effect unless and until AIP Publishing or its agent, Copyright Clearance Center, receives the Payment in accordance with Copyright Clearance Center Billing and Payment Terms and Conditions, which are incorporated herein by reference.

4. AIP Publishing or Copyright Clearance Center may, within two business days of granting this license, revoke the license for any reason whatsoever, with a full refund payable to you. Should you violate the terms of this license at any time, AIP Publishing, or Copyright Clearance Center may revoke the license with no refund to you. Notice of such revocation will be made using the contact information provided by you. Failure to receive such notice will not nullify the revocation.

5. AIP Publishing makes no representations or warranties with respect to the Material. You agree to indemnify and hold harmless AIP Publishing, and their officers, directors, employees or agents from and against any and all claims arising out of your use of the Material other than as specifically authorized herein.

6. The permission granted herein is personal to you and is not transferable or assignable without the prior written permission of AIP Publishing. This license may not be amended except in a writing signed by the party to be charged.

7. If purchase orders, acknowledgments or check endorsements are issued on any forms containing terms and conditions which are inconsistent with these provisions, such inconsistent terms and conditions shall be of no force and effect. This document, including the CCC Billing and Payment Terms and Conditions, shall be the entire agreement between the parties relating to the subject matter hereof.

This Agreement shall be governed by and construed in accordance with the laws of the State of New York. Both parties hereby submit to the jurisdiction of the courts of New York County for purposes of resolving any disputes that may arise hereunder.

V1.2

Questions? customercare@copyright.com or +1-855-239-3415 (toll free in the US) or +1-978-646-2777.

AIP PUBLISHING LICENSE
TERMS AND CONDITIONS

Mar 25, 2019

This Agreement between Mr. Jon Tong ("You") and AIP Publishing ("AIP Publishing") consists of your license details and the terms and conditions provided by AIP Publishing and Copyright Clearance Center.

License Number	4556000765625
License date	Mar 25, 2019
Licensed Content Publisher	AIP Publishing
Licensed Content Publication	Physics of Plasmas
Licensed Content Title	Three-dimensional simulations of low foot and high foot implosion experiments on the National Ignition Facility
Licensed Content Author	D. S. Clark, C. R. Weber, J. L. Milovich, et al
Licensed Content Date	May 1, 2016
Licensed Content Volume	23
Licensed Content Issue	5
Type of Use	Thesis/Dissertation
Requestor type	Student
Format	Print and electronic
Portion	Figure/Table
Number of figures/tables	2
Title of your thesis / dissertation	Ignition and Burn in Perturbed Inertial Confinement Fusion Hotspots
Expected completion date	Apr 2019
Estimated size (number of pages)	200
Requestor Location	Mr. Jon Tong Blackett Laboratory Prince Consort Road London, London SW7 2BW United Kingdom Attn: Mr. Jon Tong
Total	0.00 GBP

Terms and Conditions

AIP Publishing -- Terms and Conditions: Permissions Uses

AIP Publishing hereby grants to you the non-exclusive right and license to use and/or distribute the Material according to the use specified in your order, on a one-time basis, for the specified term, with a maximum distribution equal to the number that you have ordered. Any links or other content accompanying the Material are not the subject of this license.

1. You agree to include the following copyright and permission notice with the reproduction of the Material:"Reprinted from [FULL CITATION], with the permission of AIP Publishing." For an article, the credit line and permission notice must be printed on the first page of the article or book chapter. For photographs, covers, or tables, the notice may appear with the Material, in a footnote, or in the reference list.
2. If you have licensed reuse of a figure, photograph, cover, or table, it is your responsibility to ensure that the material is original to AIP Publishing and does not contain the copyright of another entity, and that the copyright notice of the figure, photograph, cover, or table does not indicate that it was reprinted by AIP Publishing, with permission, from another

source. Under no circumstances does AIP Publishing purport or intend to grant permission to reuse material to which it does not hold appropriate rights.

You may not alter or modify the Material in any manner. You may translate the Material into another language only if you have licensed translation rights. You may not use the Material for promotional purposes.

3. The foregoing license shall not take effect unless and until AIP Publishing or its agent, Copyright Clearance Center, receives the Payment in accordance with Copyright Clearance Center Billing and Payment Terms and Conditions, which are incorporated herein by reference.

4. AIP Publishing or Copyright Clearance Center may, within two business days of granting this license, revoke the license for any reason whatsoever, with a full refund payable to you. Should you violate the terms of this license at any time, AIP Publishing, or Copyright Clearance Center may revoke the license with no refund to you. Notice of such revocation will be made using the contact information provided by you. Failure to receive such notice will not nullify the revocation.

5. AIP Publishing makes no representations or warranties with respect to the Material. You agree to indemnify and hold harmless AIP Publishing, and their officers, directors, employees or agents from and against any and all claims arising out of your use of the Material other than as specifically authorized herein.

6. The permission granted herein is personal to you and is not transferable or assignable without the prior written permission of AIP Publishing. This license may not be amended except in a writing signed by the party to be charged.

7. If purchase orders, acknowledgments or check endorsements are issued on any forms containing terms and conditions which are inconsistent with these provisions, such inconsistent terms and conditions shall be of no force and effect. This document, including the CCC Billing and Payment Terms and Conditions, shall be the entire agreement between the parties relating to the subject matter hereof.

This Agreement shall be governed by and construed in accordance with the laws of the State of New York. Both parties hereby submit to the jurisdiction of the courts of New York County for purposes of resolving any disputes that may arise hereunder.

V1.2

Questions? customercare@copyright.com or +1-855-239-3415 (toll free in the US) or +1-978-646-2777.

Subject: Figure reuse permissions for PhD Thesis
From: "EDVARDSEN, Miriam" <M.Edvardsen@iaea.org>
Date: 26/03/2019, 12:11
To: "j.tong15@imperial.ac.uk" <j.tong15@imperial.ac.uk>

Dear Jon Tong,

Following your request for permission to reuse material:

Figure 4, from "Modeling and projecting implosion performance for the National Ignition Facility", Nucl. Fusion 59, No. 3, 032008 (2019), by D.S. Clark et al.

The IAEA is pleased to grant permission for the Figure listed above to be reproduced to the extent and for the purposes detailed in your email. Could you please also notify the corresponding author.

Concerning your request to reproduce Figure 2, please see the note in the figure caption (Adapted with permission from [26]. Copyright 2016, AIP Publishing LLC). Please contact the copyright owner directly to get their permission.

Best regards,

Ms Miriam EDVARDSEN | Publications Assistant (Marketing and Sales) |
Publishing section | Division of Conference and Document Services | Department of Management |
International Atomic Energy Agency | Vienna International Centre, PO Box 100, 1400 Vienna, Austria |
Email: **M.Edvardsen@iaea.org** | T: (+43-1) 2600-22530 | F : (+43-1) 2600-22529 |

Follow us on www.iaea.org

ATOMS FOR PEACE AND DEVELOPMENT

The IAEA Publishing Section operates an environmental management system certified to ISO 14001:2015

From: Jon Tong <j.tong15@imperial.ac.uk>
Sent: Monday, 25 March 2019 17:34
To: Sales Publication <Sales.Publications@iaea.org>
Subject: Figure reuse permissions for PhD Thesis

Dear IAEA Sales and Marketing,

I would like to request permission to reuse the figures detailed below in my PhD thesis at Imperial College London, entitled "Ignition and Burn in Perturbed Inertial Confinement Fusion hotspots".

Full title of the IAEA publication: Nuclear Fusion
Publication year: 2018, Volume 59, Issue 3
ISBN
URL : https://iopscience.iop.org/article/10.1088/1741-4326/aabcf7/meta
Exact material (pages, figures) to be reproduced: Figure 2, 4

Name: Jon Tong

Address: Blackett Laboratory, Prince Consort Road, London, SW7 2BW, United Kingdom
Email: j.tong15@imperial.ac.uk
Title of work in which it is proposed that the aforementioned IAEA material will be reproduced:
Ignition and Burn in Perturbed Inertial Confinement Fusion Hotspots
Planned publication date: April 2019
Other details of work: PhD Thesis at Imperial College London

Yours sincerely,

Jon Tong

Copyright
Clearance
Center

Confirmation Number: 11801922
Order Date: 03/26/2019

Customer Information

Customer: Jon Tong
Account Number: 3001425062
Organization: Jon Tong
Email: j.tong15@imperial.ac.uk
Phone: +44 7881835960
Payment Method: Invoice

This is not an invoice

Order Details

Plasma Physics and Controlled Fusion

Billing Status:
N/A

Order detail ID:	71861127
ISSN:	0741-3335
Publication Type:	Journal
Volume:	
Issue:	
Start page:	
Publisher:	IOP Publishing
Author/Editor:	Institute of Physics (Great Britain)

Permission Status: ✓ Granted

Permission type: Republish or display content
Type of use: Thesis/Dissertation
Order License Id: 4556631447113

Requestor type	Academic institution
Format	Print, Electronic
Portion	chart/graph/table/figure
Number of charts/graphs/tables/figures	1
The requesting person/organization	Jon Tong
Title or numeric reference of the portion(s)	Figure 6
Title of the article or chapter the portion is from	Converging geometry Rayleigh–Taylor instability and central ignition of inertial confinement fusion targets
Editor of portion(s)	N/A
Author of portion(s)	S. Atzeni, A. Schiavi, M. Temporal
Volume of serial or monograph	46
Issue, if republishing an article from a serial	12B
Page range of portion	116
Publication date of portion	2004
Rights for	Main product
Duration of use	Life of current and all future editions
Creation of copies for the disabled	no
With minor editing privileges	no
For distribution to	Worldwide
In the following language(s)	Original language of publication
	no

With incidental promotional use	
Lifetime unit quantity of new product	Up to 499
Title	Ignition and Burn in Perturbed Inertial Confinement Fusion Hotspots
Institution name	n/a
Expected presentation date	Apr 2019

Note: This item was invoiced separately through our **RightsLink service.** More info   **$ 0.00**

Total order items: 1   **Order Total: $0.00**

About Us | Privacy Policy | Terms & Conditions | Pay an Invoice

AIP PUBLISHING LICENSE
TERMS AND CONDITIONS

Mar 26, 2019

This Agreement between Mr. Jon Tong ("You") and AIP Publishing ("AIP Publishing") consists of your license details and the terms and conditions provided by AIP Publishing and Copyright Clearance Center.

License Number	4556660232297
License date	Mar 26, 2019
Licensed Content Publisher	AIP Publishing
Licensed Content Publication	Physics of Plasmas
Licensed Content Title	Nonlinear evolution of localized perturbations in the deceleration-phase Rayleigh-Taylor instability of an inertial confinement fusion capsule
Licensed Content Author	A. Schiavi, S. Atzeni
Licensed Content Date	Jul 1, 2007
Licensed Content Volume	14
Licensed Content Issue	7
Type of Use	Thesis/Dissertation
Requestor type	Student
Format	Print and electronic
Portion	Figure/Table
Number of figures/tables	1
Title of your thesis / dissertation	Ignition and Burn in Perturbed Inertial Confinement Fusion Hotspots
Expected completion date	Apr 2019
Estimated size (number of pages)	200
Requestor Location	Mr. Jon Tong Blackett Laboratory Prince Consort Road London, London SW7 2BW United Kingdom Attn: Mr. Jon Tong
Total	0.00 GBP

Terms and Conditions

AIP Publishing -- Terms and Conditions: Permissions Uses

AIP Publishing hereby grants to you the non-exclusive right and license to use and/or distribute the Material according to the use specified in your order, on a one-time basis, for the specified term, with a maximum distribution equal to the number that you have ordered. Any links or other content accompanying the Material are not the subject of this license.

1. You agree to include the following copyright and permission notice with the reproduction of the Material:"Reprinted from [FULL CITATION], with the permission of AIP Publishing." For an article, the credit line and permission notice must be printed on the first page of the article or book chapter. For photographs, covers, or tables, the notice may appear with the Material, in a footnote, or in the reference list.

2. If you have licensed reuse of a figure, photograph, cover, or table, it is your responsibility to ensure that the material is original to AIP Publishing and does not contain the copyright of another entity, and that the copyright notice of the figure, photograph, cover, or table

does not indicate that it was reprinted by AIP Publishing, with permission, from another source. Under no circumstances does AIP Publishing purport or intend to grant permission to reuse material to which it does not hold appropriate rights.

You may not alter or modify the Material in any manner. You may translate the Material into another language only if you have licensed translation rights. You may not use the Material for promotional purposes.

3. The foregoing license shall not take effect unless and until AIP Publishing or its agent, Copyright Clearance Center, receives the Payment in accordance with Copyright Clearance Center Billing and Payment Terms and Conditions, which are incorporated herein by reference.

4. AIP Publishing or Copyright Clearance Center may, within two business days of granting this license, revoke the license for any reason whatsoever, with a full refund payable to you. Should you violate the terms of this license at any time, AIP Publishing, or Copyright Clearance Center may revoke the license with no refund to you. Notice of such revocation will be made using the contact information provided by you. Failure to receive such notice will not nullify the revocation.

5. AIP Publishing makes no representations or warranties with respect to the Material. You agree to indemnify and hold harmless AIP Publishing, and their officers, directors, employees or agents from and against any and all claims arising out of your use of the Material other than as specifically authorized herein.

6. The permission granted herein is personal to you and is not transferable or assignable without the prior written permission of AIP Publishing. This license may not be amended except in a writing signed by the party to be charged.

7. If purchase orders, acknowledgments or check endorsements are issued on any forms containing terms and conditions which are inconsistent with these provisions, such inconsistent terms and conditions shall be of no force and effect. This document, including the CCC Billing and Payment Terms and Conditions, shall be the entire agreement between the parties relating to the subject matter hereof.

This Agreement shall be governed by and construed in accordance with the laws of the State of New York. Both parties hereby submit to the jurisdiction of the courts of New York County for purposes of resolving any disputes that may arise hereunder.

V1.2

Questions? customercare@copyright.com or +1-855-239-3415 (toll free in the US) or +1-978-646-2777.

AIP PUBLISHING LICENSE
TERMS AND CONDITIONS

Apr 09, 2019

This Agreement between Mr. Jon Tong ("You") and AIP Publishing ("AIP Publishing") consists of your license details and the terms and conditions provided by AIP Publishing and Copyright Clearance Center.

License Number	4564710411138
License date	Apr 09, 2019
Licensed Content Publisher	AIP Publishing
Licensed Content Publication	Physics of Plasmas
Licensed Content Title	Development of the indirect-drive approach to inertial confinement fusion and the target physics basis for ignition and gain
Licensed Content Author	John Lindl
Licensed Content Date	Nov 1, 1995
Licensed Content Volume	2
Licensed Content Issue	11
Type of Use	Thesis/Dissertation
Requestor type	Student
Format	Print and electronic
Portion	Figure/Table
Number of figures/tables	1
Title of your thesis / dissertation	Ignition and Burn in Perturbed Inertial Confinement Fusion Hotspots
Expected completion date	Apr 2019
Estimated size (number of pages)	200
Requestor Location	Mr. Jon Tong Blackett Laboratory Prince Consort Road London, London SW7 2BW United Kingdom Attn: Mr. Jon Tong
Total	**0.00 GBP**
Terms and Conditions	

AIP Publishing -- Terms and Conditions: Permissions Uses

AIP Publishing hereby grants to you the non-exclusive right and license to use and/or distribute the Material according to the use specified in your order, on a one-time basis, for the specified term, with a maximum distribution equal to the number that you have ordered. Any links or other content accompanying the Material are not the subject of this license.

1. You agree to include the following copyright and permission notice with the reproduction of the Material:"Reprinted from [FULL CITATION], with the permission of AIP Publishing." For an article, the credit line and permission notice must be printed on the first page of the article or book chapter. For photographs, covers, or tables, the notice may appear with the Material, in a footnote, or in the reference list.

2. If you have licensed reuse of a figure, photograph, cover, or table, it is your responsibility to ensure that the material is original to AIP Publishing and does not contain the copyright of another entity, and that the copyright notice of the figure, photograph, cover, or table does not indicate that it was reprinted by AIP Publishing, with permission, from another source. Under no circumstances does AIP Publishing purport or intend to grant permission to reuse material to which it does not hold appropriate rights.

   You may not alter or modify the Material in any manner. You may translate the Material into another language only if you have licensed translation rights. You may not use the Material for promotional purposes.

3. The foregoing license shall not take effect unless and until AIP Publishing or its agent, Copyright Clearance Center, receives the Payment in accordance with Copyright Clearance Center Billing and Payment Terms and Conditions, which are incorporated herein by reference.

4. AIP Publishing or Copyright Clearance Center may, within two business days of granting this license, revoke the license for any reason whatsoever, with a full refund payable to you. Should you violate the terms of this license at any time, AIP Publishing, or Copyright Clearance Center may revoke the license with no refund to you. Notice of such revocation will be made using the contact information provided by you. Failure to receive such notice will not nullify the revocation.

5. AIP Publishing makes no representations or warranties with respect to the Material. You agree to indemnify and hold harmless AIP Publishing, and their officers, directors, employees or agents from and against any and all claims arising out of your use of the Material other than as specifically authorized herein.

6. The permission granted herein is personal to you and is not transferable or assignable without the prior written permission of AIP Publishing. This license may not be amended except in a writing signed by the party to be charged.

7. If purchase orders, acknowledgments or check endorsements are issued on any forms containing terms and conditions which are inconsistent with these provisions, such inconsistent terms and conditions shall be of no force and effect. This document, including the CCC Billing and Payment Terms and Conditions, shall be the entire agreement between the parties relating to the subject matter hereof.

This Agreement shall be governed by and construed in accordance with the laws of the State of New York. Both parties hereby submit to the jurisdiction of the courts of New York County for purposes of resolving any disputes that may arise hereunder.

V1.2

Questions? customercare@copyright.com or +1-855-239-3415 (toll free in the US) or +1-978-646-2777.

PARTIES:

1. **Oxford Publishing Limited** (Company number – 01748118) (Licensor); and
2. **Jon Tong** (Licensee).

Thank you for your recent permission request. Some permission requests for use of material published by the Licensor, such as this one, are now being facilitated by PLSclear.

Set out in this licence cover sheet (the **Licence Cover Sheet**) are the principal commercial terms under which Licensor has agreed to license certain Licensed Material (as defined below) to Licensee. The terms in this Licence Cover Sheet are subject to the attached General Terms and Conditions, which together with this Licence Cover Sheet constitute the licence agreement (the **Licence**) between Licensor and Licensee as regards the Licensed Material. The terms set out in this Licence Cover Sheet take precedence over any conflicting provision in the General Terms and Conditions.

Licence Terms

Licence Date:	09/04/2019
PLSclear Ref No:	12917

The Licensor

Company name:	Oxford Publishing Limited
Address:	Rights Department
	Great Clarendon Street
	Oxford
	OX2 6DP
	GB

The Licensee

Licensee Contact Name:	Jon Tong
Licensee Address:	Blackett Laboratory
	Prince Consort Road
	London
	SW7 2BW
	United Kingdom

Licensed Material

title:	The Physics of Inertial Fusion BeamPlasma Interaction, Hydrodynamics, Hot Dense Matter
ISBN/ISSN:	9780198562641
publisher:	Oxford Publishing Limited

figure number & title / caption	4.10 The 1D simulation of ignition and burn of an initially isobaric, equimolar DT configuration. Sequences of radial profiles of ion temperature (a), density (b) and pressure (c) at selected times.
page number	89
position on page	bottom-right
reproduction colour	Black and White
reproduction size	Full page
positioning	inside or later pages
Are you requesting permission to reuse your own work?	Yes. I am the author
Are you using the content as a prop?	content will NOT be used as a prop

For Use In Licensee's Publication(s)

usage type	Book, Journal, Magazine or Academic Paper...-Thesis
estimated publication date	April 2019
language	English
number of pages	203
other relevant Information	Reproducing figure for PhD thesis work.
publication title	Ignition and Burn in Perturbed Inertial Confinement Fusion Hotspots
type of document	PhD Thesis

Rights Granted

Exclusivity:	Non-Exclusive
Format:	THESIS/WHITEPAPER/CONULTATION DOCUMENT
Language:	English
Territory:	UK & Commonwealth
Duration:	Lifetime of Licensee's edition
Maximum Circulation:	Total: 1 copies
Additional Terms:	If at some future date your thesis is published it will be necessary to re-clear this permission. Please also note that if the material to be used is acknowledged to any other source, you will need to clear permission with the rights holder and for any electronic version the © line must appear on the same page as the OUP material and the OUP material should not be included under a Creative Commons license, or any other open-access license allowing onward reuse.

Payment Details

Fee Payable:	£0.00 [+ VAT if applicable]

Payment Terms: Strictly 30 days from date of Licence

GENERAL TERMS AND CONDITIONS

1. Definitions and Interpretation

1.1 Capitalised words and expressions in these General Terms and Conditions have the meanings given to them in the Licence Cover Sheet.

1.2 In this Licence any references (express or implied) to statutes or provisions are references to those statutes or provisions as amended or re-enacted from time to time. The term **including** will be construed as illustrative, without limiting the sense or scope of the words preceding it. A reference to in **writing** or **written** includes faxes and email. The singular includes the plural and vice versa.

2. Grant of Rights

2.1 Subject to payment by Licensee of the Licence Fee in accordance with paragraph 3 below, Licensor grants to Licensee the non-exclusive right to use the Licensed Material as specified in the Licence Cover Sheet.

2.2 The rights licensed to Licensee under this Licence do not include the right to use any third party copyright material incorporated in the Licensed Material. Licensee should check the Licensed Material carefully and seek permission for the use of any such third party copyright material from the relevant copyright owner(s).

2.3 Unless otherwise stated in the Licence Cover Sheet, the Licensed Material may be:

2.3.1 subjected to minor editing, including for the purposes of creating alternative formats to provide access for a beneficiary person (provided that any such editing does not amount to derogatory treatment); and/or

2.3.2 used for incidental promotional use (such as online retail providers' search facilities).

2.4 Save as expressly permitted in this Licence or as otherwise permitted by law, no use or modification of the Licensed Material may be made by Licensee without Licensor's prior written permission.

3. Payment

3.1 Licensee must pay to Licensor the Licence Fee by means of either credit card or on receipt of an invoice, as selected by Licensee during the licence application process via the PLSclear service.

3.2 If payment is by invoice, Licensee agrees to pay the Licence Fee in full by no later than the payment date specified in the relevant invoice.

4. Copyright Notice and Acknowledgement

4.1 Licensee must ensure that the following notices and acknowledgements are reproduced prominently alongside each reproduction by Licensee of the Licensed Material:

4.1.1 the title and author of the Licensed Material;

4.1.2 the copyright notice included in the Licensed Material; and

4.1.3 the statement "Reproduced with permission of the Licensor through PLSclear."

5. **Reversion of Rights**

5.1 The rights licensed to Licensee under this Licence will terminate immediately and automatically upon the earliest of the following events to occur:

5.1.1 the Licence Fee not being received by Licensor in full by the payment date specified in the relevant invoice;

5.1.2 the Licensed Material not being used by Licensee within 18 months of the Licence Date;

5.1.3 expiry of the Licence Duration; or

5.1.4 the Maximum Circulation being reached.

6. **Miscellaneous**

6.1 By using the Licensed Material, Licensee will be deemed to have accepted all the terms and conditions contained in this Licence.

6.2 This Licence contains the entire understanding and agreement of the parties relating to its subject matter and supersedes in all respects any previous or other existing arrangements, agreements or understandings between the parties whether oral or written in relation to its subject matter.

6.3 Licensee may not assign this Licence or any of its rights or obligations hereunder to any third party without Licensor's prior written consent.

6.4 This Licence is governed by and shall be construed in accordance with the laws of England and Wales and the parties hereby irrevocably submit to the non-exclusive jurisdiction of the Courts of England and Wales as regards any claim, dispute or matter arising under or in relation to this Licence.

AIP PUBLISHING LICENSE
TERMS AND CONDITIONS

Jun 16, 2019

This Agreement between Mr. Jon Tong ("You") and AIP Publishing ("AIP Publishing") consists of your license details and the terms and conditions provided by AIP Publishing and Copyright Clearance Center.

License Number	4610870056028
License date	Jun 16, 2019
Licensed Content Publisher	AIP Publishing
Licensed Content Publication	Physics of Plasmas
Licensed Content Title	Diagnostic signatures of performance degrading perturbations in inertial confinement fusion implosions
Licensed Content Author	K. McGlinchey, B. D. Appelbe, A. J. Crilly, et al
Licensed Content Date	Dec 1, 2018
Licensed Content Volume	25
Licensed Content Issue	12
Type of Use	Thesis/Dissertation
Requestor type	Author (original article)
Format	Print and electronic
Portion	Figure/Table
Number of figures/tables	1
Title of your thesis / dissertation	Ignition and Burn in Perturbed Inertial Confinement Fusion Hotspots
Expected completion date	Apr 2019
Estimated size (number of pages)	200
Requestor Location	Mr. Jon Tong Blackett Laboratory Prince Consort Road London, London SW7 2BW United Kingdom Attn: Mr. Jon Tong
Total	0.00 USD

Terms and Conditions

source. Under no circumstances does AIP Publishing purport or intend to grant permission to reuse material to which it does not hold appropriate rights.

You may not alter or modify the Material in any manner. You may translate the Material into another language only if you have licensed translation rights. You may not use the Material for promotional purposes.

3. The foregoing license shall not take effect unless and until AIP Publishing or its agent, Copyright Clearance Center, receives the Payment in accordance with Copyright Clearance Center Billing and Payment Terms and Conditions, which are incorporated herein by reference.

4. AIP Publishing or Copyright Clearance Center may, within two business days of granting this license, revoke the license for any reason whatsoever, with a full refund payable to you. Should you violate the terms of this license at any time, AIP Publishing, or Copyright Clearance Center may revoke the license with no refund to you. Notice of such revocation will be made using the contact information provided by you. Failure to receive such notice will not nullify the revocation.

5. AIP Publishing makes no representations or warranties with respect to the Material. You agree to indemnify and hold harmless AIP Publishing, and their officers, directors, employees or agents from and against any and all claims arising out of your use of the Material other than as specifically authorized herein.

6. The permission granted herein is personal to you and is not transferable or assignable without the prior written permission of AIP Publishing. This license may not be amended except in a writing signed by the party to be charged.

7. If purchase orders, acknowledgments or check endorsements are issued on any forms containing terms and conditions which are inconsistent with these provisions, such inconsistent terms and conditions shall be of no force and effect. This document, including the CCC Billing and Payment Terms and Conditions, shall be the entire agreement between the parties relating to the subject matter hereof.

This Agreement shall be governed by and construed in accordance with the laws of the State of New York. Both parties hereby submit to the jurisdiction of the courts of New York County for purposes of resolving any disputes that may arise hereunder.

V1.2

Questions? customercare@copyright.com or +1-855-239-3415 (toll free in the US) or +1-978-646-2777.

Bibliography

[1] Edwards, M. J., Patel, P. K., Lindl, J. D., Atherton, L. J., Glenzer, S. H., Haan, S. W., Kilkenny, J. D., Landen, O. L., Moses, E. I., Nikroo, A., Petrasso, R., Sangster, T. C., Springer, P. T., Batha, S., Benedetti, R., Bernstein, L., Betti, R., Bleuel, D. L., Boehly, T. R., Bradley, D. K., Caggiano, J. A., Callahan, D. A., Celliers, P. M., Cerjan, C. J., Chen, K. C., Clark, D. S., Collins, G. W., Dewald, E. L., Divol, L., Dixit, S., Doeppner, T., Edgell, D. H., Fair, J. E., Farrell, M., Fortner, R. J., Frenje, J., Gatu Johnson, M. G., Giraldez, E., Glebov, V. Y., Grim, G., Hammel, B. A., Hamza, A. V., Harding, D. R., Hatchett, S. P., Hein, N., Herrmann, H. W., Hicks, D., Hinkel, D. E., Hoppe, M., Hsing, W. W., Izumi, N., Jacoby, B., Jones, O. S., Kalantar, D., Kauffman, R., Kline, J. L., Knauer, J. P., Koch, J. A., Kozioziemski, B. J., Kyrala, G., LaFortune, K. N., Le Pape, S., Leeper, R. J., Lerche, R., Ma, T., MacGowan, B. J., MacKinnon, A. J., Macphee, A., Mapoles, E. R., Marinak, M. M., Mauldin, M., McKenty, P. W., Meezan, M., Michel, P. A., Milovich, J., Moody, J. D., Moran, M., Munro, D. H., Olson, C. L., Opachich, K., Pak, A. E., Parham, T., Park, H.-S., Ralph, J. E., Regan, S. P., Remington, B., Rinderknecht, H., Robey, H. F., Rosen, M., Ross, S., Salmonson, J. D., Sater, J., Schneider, D. H., Séguin, F. H., Sepke, S. M., Shaughnessy, D. A., Smalyuk, V. A., Spears, B. K., Stoeckl, C., Stoeffl, W., Suter, L., Thomas, C. A., Tommasini, R., Town, R. P., Weber, S. V., Wegner, P. J., Widman, K., Wilke, M., Wilson, D. C., Yeamans, C. B., and Zylstra, A. *Phys. Plasmas* **20**(7), 070501 (2013). 15, 27, 180

[2] Clark, D. S., Weber, C. R., Milovich, J. L., Salmonson, J. D., Kritcher, A. L., Haan, S. W., Hammel, B. A., Hinkel, D. E., Hurricane, O. A., Jones, O. S., Marinak, M. M., Patel, P. K., Robey, H. F., Sepke, S. M., and Edwards, M. J. *Phys. Plasmas* **23**(5), 056302 (2016). 15, 32, 33, 34, 141, 143, 146, 180

[3] Clark, D. S., Weber, C. R., Kritcher, A. L., Milovich, J. L., Patel, P. K., Haan, S. W., Hammel, B. A., Koning, J. M., Marinak, M. M., Patel, M. V., Schroeder, C. R., Sepke, S. M., and Edwards, M. J. *Nucl. Fusion* **59**, 032008 (2019). 15, 22, 33, 34, 40, 59, 141, 143, 150, 164, 165, 167, 171, 180

[4] Atzeni, S., Schiavi, A., and Temporal, M. *Plasma Phys. Control. Fusion* **46**(46), 111–120 (2004). 15, 37, 38, 180

[5] Schiavi, A. and Atzeni, S. *Phys. Plasmas* **14**(7), 1–5 (2007). 15, 37, 39, 180

[6] Lindl, J. *Phys. Plasmas* **2**(11), 3933–4024 (1995). 15, 31, 49, 180

[7] Atzeni, S. and Meyer-ter Vehn, J. *The Physics of Inertial Fusion: Beam Plasma Interaction, Hydrodynamics, Hot Dense Matter*. International Series of Monographs on Physics. OUP Oxford, (2004). 16, 26, 29, 56, 84, 85, 100, 180

[8] McGlinchey, K., Appelbe, B. D., Crilly, A. J., Tong, J. K., Walsh, C. A., and Chittenden, J. P. *Phys. Plasmas* **25**(12), 122705 (2018). 20, 54, 96, 143, 155, 180

[9] Le Pape, S., Hopkins, L. F. B., Divol, L., Pak, A., Dewald, E. L., Bhandarkar, S., Bennedetti, L. R., Bunn, T., Biener, J., Crippen, J., Casey, D., Edgell, D., Fittinghoff, D. N., Goyon, C., Haan, S., Hatarik, R., Havre, M., Ho, D. D.-m., Izumi, N., Jaquez, J., Khan, S. F., Kyrala, G. A., Ma, T., Mackinnon, A. J., Macphee, A. G., Macgowan, B. J., Meezan, N. B., Milovich, J., Millot, M., Michel, P., Nagel, S. R., Nikroo, A., Patel, P., Ralph, J., Ross, J. S., Rice, N. G., Strozzi, D., Stadermann, M., Volegov, P., Yeamans, C., Weber, C., Wild, C., Callahan, D., and Hurricane, O. A. *Phys. Rev. Lett.* **120**(24), 245003 (2018). 22, 34, 37, 141, 150, 164, 165, 167

[10] UN. *Paris Agreement* (2015). 24

[11] Gür, T. M. *Energy Environ. Sci.* **11**(10), 2696–2767 (2018). 24

[12] Fraley, G. S., Linnebur, E. J., Mason, R. J., and Morse, R. L. *Phys. Fluids* **17**(1974), 474 (1974). 26, 44, 52, 63

[13] Chu, M. S. *Phys. Fluids* **15**(3), 413–422 (1972). 26

[14] Boehly, T., Brown, D., Craxton, R., Keck, R., Knauer, J., Kelly, J., Kessler, T., Kumpan, S., Loucks, S., Letzring, S., Marshall, F., McCrory, R., Morse, S., Seka, W., Soures, J., and Verdon, C. *Opt. Commun.* **133**(1-6), 495–506 (1997). 29, 58

[15] Obenschain, S. P., Bodner, S. E., Colombant, D., Gerber, K., Lehmberg, R. H., McLean, E. a., Mostovych, a. N., Pronko, M. S., Pawley, C. J., Schmitt, a. J., Sethian, J. D., Serlin, V., Stamper, J. a., Sullivan, C. a., Dahlburg, J. P., Gardner, J. H., Chan, Y., Deniz, a. V., Hardgrove, J., Lehecka, T., and Klapisch, M. *Phys. Plasmas* **3**(5), 2098 (1996). 29

[16] Azechi, H., Jitsuno, T., Kanabe, T., Katayama, M., Mima, K., Miyanaga, N., Nakai, M., Nakai, S., Nakaishi, H., Nakatsuka, M., Nishiguchi, a., Norrays, P. a., Setsuhara, Y., Takagi, M., Yamanaka, M., and Yamanaka, C. *Laser Part. Beams* **9**(02), 193 (2009). 29

[17] Skupsky, S., Marozas, J. A., Craxton, R. S., Betti, R., Collins, T. J., Delettrez, J. A., Goncharov, V. N., McKenty, P. W., Radha, P. B., Boehly, T. R., Knauer, J. P., Marshall, F. J., Harding, D. R., Kilkenny, J. D., Meyerhofer, D. D., Sangster, T. C., and McCrory, R. L. *Phys. Plasmas* **11**(5 PART 2), 2763–2770 (2004). 29

[18] Hohenberger, M., Radha, P. B., Myatt, J. F., Le Pape, S., Marozas, J. A., Marshall, F. J., Michel, D. T., Regan, S. P., Seka, W., Shvydky, A., Sangster, T. C., Bates, J., Betti, R., Boehly, T. R., Bonino, M. J., Casey, D. T., Collins, T. J., Craxton, R. S., Delettrez, J. A., Edgell, D. H., Epstein, R., Fiksel, G., Fitzsimmons, P., Frenje, J. A., Froula, D. H., Goncharov, V. N., Harding, D. R., Hu, S. X., Kalantar, D. H., Karasik, M., Kessler, T. J., Kilkenny, J. D., Knauer, J. P., Kurz, C., Lafon, M., La Fortune, K. N., MacGowan, B. J., Mackinnon, A. J., MacPhee, A. G., McCrory, R. L., McKenty, P. W., Meeker, J. F., Meyerhofer, D. D., Nagel, S. R., Nikroo, A., Obenschein, S., Petrasso, R. D., Ralph, J. E., Rinderknecht, H. G., Rosenberg, M. J., Schmitt, A. J., Wallace, R. J., Weaver, J., Widmayer, C. C., Skupsky, S., Solodov, A. A., Stoeckl, C., Yaakobi, B., Zuegel, J. D., and Zylstra, A. *Phys. Plasmas* **22**(056308) (2015). 29

[19] Tabak, M., Hammer, J., Glinsky, M. E., Kruer, W. L., Wilks, S. C., Woodworth, J., Campbell, E. M., Perry, M. D., and Mason, R. J. *Phys. Plasmas* **1**(5), 1626 (1994). 30

[20] Roth, M., Cowan, T. E., Key, M. H., Hatchett, S. P., Brown, C., Fountain, W., Johnson, J., Pennington, D. M., Snavely, R. A., Wilks, S. C., Yasuike, K., Ruhl, H., Pegoraro, F., Bulanov, S. V., Campbell, E. M., Perry, M. D., and Powell, H. *Phys. Rev. Lett.* **86**(3), 436–439 (2001). 30

[21] Wilks, S. C., Kruer, W. L., Tabak, M., and Langdon, A. B. *Phys. Rev. Lett.* **69**(9), 1383–1386 (1992). 30

[22] Norreys, P. A., Allott, R., Clarke, R. J., Collier, J., Neely, D., Rose, S. J., Zepf, M., Santala, M., Bell, A. R., Krushelnick, K., Dangor, A. E., Woolsey, N. C., Evans, R. G., Habara, H., Norimatsu, T., and Kodama, R. *Phys. Plasmas* **7**(9), 3721–3726 (2000). 30

[23] Betti, R., Zhou, C. D., Anderson, K. S., Perkins, L. J., Theobald, W., and Solodov, A. A. *Phys. Rev. Lett.* **98**(15) (2007). 30

[24] Perkins, L. J., Betti, R., Lafortune, K. N., and Williams, W. H. *Phys. Rev. Lett.* **103**, 045004 (2009). 30

[25] Atzeni, S., Ribeyre, X., Schurtz, G., Schmitt, A., Canaud, B., Betti, R., and Perkins, L. *Nucl. Fusion* **54**(5), 054008 (2014). 31

[26] Gotchev, O. V., Chang, P. Y., Knauer, J. P., Meyerhofer, D. D., Polomarov, O., Frenje, J., Li, C. K., Manuel, M. J. E., Petrasso, R. D., Rygg, J. R., Séguin, F. H., and Betti, R. *Phys. Rev. Lett.* **103**, 215004 (2009). 31

[27] Perkins, L. J., Logan, B. G., Zimmerman, G. B., and Werner, C. J. *Phys. Plasmas* **20**, 072708 (2013). 31, 173

[28] Slutz, S. A., Herrmann, M. C., Vesey, R. A., Sefkow, A. B., Sinars, D. B., Rovang, D. C., Peterson, K. J., and Cuneo, M. E. *Phys. Plasmas* **17**, 056303 (2010). 31

[29] Gomez, M. R., Slutz, S. A., Sefkow, A. B., Sinars, D. B., Hahn, K. D., Hansen, S. B., Harding, E. C., Knapp, P. F., Schmit, P. F., Jennings, C. A., Awe, T. J., Geissel, M., Rovang, D. C., Chandler, G. A., Cooper, G. W., Cuneo, M. E., Harvey-Thompson, A. J., Herrmann, M. C., Hess, M. H., Johns, O., Lamppa, D. C., Martin, M. R., McBride, R. D., Peterson, K. J., Porter, J. L., Robertson, G. K., Rochau, G. A., Ruiz, C. L., Savage, M. E., Smith, I. C., Stygar, W. A., and Vesey, R. A. *Phys. Rev. Lett.* **113**, 155003 (2014). 31

[30] Amendt, P., Colvin, J. D., Tipton, R. E., Hinkel, D. E., Edwards, M. J., Landen, O. L., Ramshaw, J. D., Suter, L. J., Varnum, W. S., and Watt, R. G. *Phys. Plasmas* **9**(5), 2221–2233 (2002). 31

[31] Amendt, P., Cerjan, C., Hamza, A., Hinkel, D. E., Milovich, J. L., and Robey, H. F. *Phys. Plasmas* **14**, 056312 (2007). 31

[32] Molvig, K., Schmitt, M. J., Betti, R., Michael Campbell, E., and McKenty, P. *Phys. Plasmas* **25**, 082708 (2018). 31

[33] Lindl, J. D., Amendt, P., Berger, R. L., Glendinning, S. G., Glenzer, S. H., Haan, S. W., Kauffman, R. L., Landen, O. L., and Suter, L. J. *Phys. Plasmas* **11**(2), 339 (2004). 31

[34] Lindl, J., Landen, O., Edwards, J., and Moses, E. *Phys. Plasmas* **21**, 020501 (2014). 32

[35] Scott, R. H. H., Clark, D. S., Bradley, D. K., Callahan, D. A., Edwards, M. J., Haan, S. W., Jones, O. S., Spears, B. K., Marinak, M. M., Town, R. P. J., Norreys, P. A., and Suter, L. J. *Phys. Rev. Lett.* **110**, 075001 (2013). 32

[36] Hurricane, O. A., Callahan, D. A., Casey, D. T., Dewald, E. L., Dittrich, T. R., Doeppner, T., Barrios Garcia, M. A., Hinkel, D. E., Berzak Hopkins, L. F., Kervin, P., Kline, J. L., Le Pape, S., Ma, T., Macphee, A. G., Milovich, J. L., Moody, J., Pak, A. E., Patel, P. K., Park, H. S., Remington, B. A., Robey, H. F., Salmonson, J. D., Springer, P. T., Tommasini, R., Benedetti, L. R., Caggiano, J. A., Celliers, P., Cerjan, C., Dylla-Spears, R., Edgell, D., Edwards, M. J., Fittinghoff, D., Grim, G. P., Guler, N., Izumi, N., Frenje, J. A., Gatu Johnson, M., Haan, S., Hatarik, R., Herrmann, H., Khan, S., Knauer, J., Kozioziemski, B. J., Kritcher, A. L., Kyrala, G., Maclaren, S. A., Merrill, F. E., Michel, P., Ralph, J., Ross, J. S., Rygg, J. R., Schneider, M. B., Spears, B. K., Widmann, K., and Yeamans, C. B. *Phys. Plasmas* **21**, 056314 (2014). 32, 141

[37] Park, H.-S., Hurricane, O. A., Callahan, D. A., Casey, D. T., Dewald, E. L., Dittrich, T. R., Döppner, T., Hinkel, D. E., Berzak Hopkins, L. F., Le Pape, S.,

Ma, T., Patel, P. K., Remington, B. A., Robey, H. F., Salmonson, J. D., and Kline, J. L. *Phys. Rev. Lett.* **112**(5), 055001 (2014). 32

[38] Dittrich, T. R., Hurricane, O. A., Callahan, D. A., Dewald, E. L., Döppner, T., Hinkel, D. E., Berzak Hopkins, L. F., Le Pape, S., Ma, T., Milovich, J. L., Moreno, J. C., Patel, P. K., Park, H.-S., Remington, B. A., Salmonson, J. D., and Kline, J. L. *Phys. Rev. Lett.* **112**(5), 055002 (2014). 32

[39] Hurricane, O. A., Callahan, D. A., Casey, D. T., Dewald, E. L., Dittrich, T. R., Döppner, T., Haan, S., Hinkel, D. E., Berzak Hopkins, L. F., Jones, O., Kritcher, A. L., Le Pape, S., Ma, T., MacPhee, A. G., Milovich, J. L., Moody, J., Pak, A., Park, H.-S., Patel, P. K., Ralph, J. E., Robey, H. F., Ross, J. S., Salmonson, J. D., Spears, B. K., Springer, P. T., Tommasini, R., Albert, F., Benedetti, L. R., Bionta, R., Bond, E., Bradley, D. K., Caggiano, J., Celliers, P. M., Cerjan, C., Church, J. A., Dylla-Spears, R., Edgell, D., Edwards, M. J., Fittinghoff, D., Barrios Garcia, M. A., Hamza, A., Hatarik, R., Herrmann, H., Hohenberger, M., Hoover, D., Kline, J. L., Kyrala, G., Kozioziemski, B., Grim, G., Field, J. E., Frenje, J., Izumi, N., Gatu Johnson, M., Khan, S. F., Knauer, J., Kohut, T., Landen, O., Merrill, F., Michel, P., Moore, A., Nagel, S. R., Nikroo, A., Parham, T., Rygg, R. R., Sayre, D., Schneider, M., Shaughnessy, D., Strozzi, D., Town, R. P. J., Turnbull, D., Volegov, P., Wan, A., Widmann, K., Wilde, C., and Yeamans, C. *Nat. Phys.* **12**(April), 800–806 (2016). 32, 133, 163

[40] Kritcher, A. L., Hinkel, D. E., Callahan, D. A., Hurricane, O. A., Clark, D. S., Casey, D. T., Dewald, E. L., Dittrich, T. R., Döppner, T., Barrios Garcia, M. A., Haan, S., Berzak Hopkins, L. F., Jones, O., Landen, O., Ma, T., Meezan, N., Milovich, J. L., Pak, A. E., Park, H. S., Patel, P. K., Ralph, J., Robey, H. F., Salmonson, J. D., Sepke, S., Spears, B., Springer, P. T., Thomas, C. A., Town, R., Celliers, P. M., and Edwards, M. J. *Phys. Plasmas* **23**, 052709 (2016). 32, 54

[41] Marinak, M. M., Kerbel, G. D., Gentile, N. A., Jones, O., Munro, D., Pollaine, S., Dittrich, T. R., and Haan, S. W. *Phys. Plasmas* **8**(5 II), 2275–2280 (2001). 32

[42] Kritcher, A. L., Town, R., Bradley, D., Clark, D. S., Spears, B., Jones, O., Haan, S., Springer, P. T., Lindl, J., Scott, R. H. H., Callahan, D., Edwards, M. J., and Landen, O. L. *Phys. Plasmas* **21**, 042708 (2014). 32

[43] Michel, P., Glenzer, S. H., Divol, L., Bradley, D. K., Callahan, D., Dixit, S., Glenn, S., Hinkel, D., Kirkwood, R. K., Kline, J. L., Kruer, W. L., Kyrala, G. A., Le Pape, S., Meezan, N. B., Town, R., Widmann, K., Williams, E. A., MacGowan, B. J., Lindl, J., and Suter, L. J. *Phys. Plasmas* **17**(5) (2010). 33

[44] Kruer, W. L. *The Physics of Laser Plasma Interactions.* Addison-Wesley, (1988). 33

[45] Ralph, J. E., Landen, O., Divol, L., Pak, A., Ma, T., Callahan, D. A., Kritcher, A. L., Döppner, T., Hinkel, D. E., Jarrott, C., Moody, J. D., Pollock, B. B., Hurricane, O., and Edwards, M. J. *Phys. Plasmas* **25**(8), 082701 (2018). 33, 36

[46] Nagel, S. R., Haan, S. W., Rygg, J. R., Barrios, M., Benedetti, L. R., Bradley, D. K., Field, J. E., Hammel, B. A., Izumi, N., Jones, O. S., Khan, S. F., Ma, T., Pak, A. E., Tommasini, R., and Town, R. P. J. *Phys. Plasmas* **22**, 022704 (2015). 33

[47] Tommasini, R., Field, J. E., Hammel, B. A., Landen, O. L., Haan, S. W., Aracne-Ruddle, C., Benedetti, L. R., Bradley, D. K., Callahan, D. A., Dewald, E. L., Doeppner, T., Edwards, M. J., Hurricane, O. A., Izumi, N., Jones, O. A., Ma, T., Meezan, N. B., Nagel, S. R., Rygg, J. R., Segraves, K. S., Stadermann, M., Strauser, R. J., and Town, R. P. J. *Phys. Plasmas* **22**, 056315 (2015). 33

[48] Smalyuk, V. A., Weber, S. V., Casey, D. T., Clark, D. S., Field, J. E., Haan, S. W., Hammel, B. A., Hamza, A. V., Hoover, D. E., Landen, O. L., Nikroo, A., Robey, H. F., and Weber, C. R. *Phys. Plasmas* **22**, 072704 (2015). 33

[49] Walsh, C. A., Chittenden, J. P., Mcglinchey, K., Niasse, N. P. L., and Appelbe, B. D. *Phys. Rev. Lett.* **118**, 155001 (2017). 33

[50] MacKinnon, A. J., Meezan, N. B., Ross, J. S., Le Pape, S., Berzak Hopkins, L., Divol, L., Ho, D., Milovich, J., Pak, A., Ralph, J., Döppner, T., Patel, P. K., Thomas, C., Tommasini, R., Haan, S., Macphee, A. G., McNaney, J., Caggiano, J., Hatarik, R., Bionta, R., Ma, T., Spears, B., Rygg, J. R., Benedetti, L. R., Town, R. P. J., Bradley, D. K., Dewald, E. L., Fittinghoff, D., Jones, O. S., Robey, H. R., Moody, J. D., Khan, S., Callahan, D. A., Hamza, A., Biener, J., Celliers, P. M., Braun, D. G., Erskine, D. J., Prisbrey, S. T., Wallace, R. J., Kozioziemski, B., Dylla-Spears, R., Sater, J., Collins, G., Storm, E., Hsing, W., Landen, O., Atherton, J. L., Lindl, J. D., Edwards, M. J., Frenje, J. A., Gatu-Johnson, M., Li, C. K., Petrasso, R., Rinderknecht, H., Rosenberg, M., Séguin, F. H., Zylstra, A., Knauer, J. P., Grim, G., Guler, N., Merrill, F., Olson, R., Kyrala, G. A., Kilkenny, J. D., Nikroo, A., Moreno, K., Hoover, D. E., Wild, C., and Werner, E. *Phys. Plasmas* **21**, 056318 (2014). 34

[51] Le Pape, S., Berzak Hopkins, L. F., Divol, L., Meezan, N., Turnbull, D., MacKinnon, A. J., Ho, D., Ross, J. S., Khan, S., Pak, A., Dewald, E., Benedetti, L. R., Nagel, S., Biener, J., Callahan, D. A., Yeamans, C., Michel, P., Schneider, M., Kozioziemski, B., Ma, T., MacPhee, A. G., Haan, S., Izumi, N., Hatarik, R., Sterne, P., Celliers, P., Ralph, J., Rygg, R., Strozzi, D., Kilkenny, J., Rosenberg, M., Rinderknecht, H., Sio, H., Gatu-Johnson, M., Frenje, J., Petrasso, R., Zylstra, A., Town, R., Hurricane, O., Nikroo, A., and Edwards, M. J. *Phys. Plasmas* **23**, 056311 (2016). 34

[52] Meezan, N. B., Berzak Hopkins, L. F., Le Pape, S., Divol, L., Mackinnon, A. J., Döppner, T., Ho, D. D., Jones, O. S., Khan, S. F., Ma, T., Milovich, J. L., Pak, A. E., Ross, J. S., Thomas, C. A., Benedetti, L. R., Bradley, D. K., Celliers, P. M., Clark, D. S., Field, J. E., Haan, S. W., Izumi, N., Kyrala, G. A., Moody, J. D., Patel, P. K., Ralph, J. E., Rygg, J. R., Sepke, S. M., Spears, B. K., Tommasini, R., Town, R. P. J., Biener, J., Bionta, R. M., Bond, E. J., Caggiano, J. A., Eckart, M. J., Gatu Johnson, M., Grim, G. P., Hamza, A. V., Hartouni, E. P., Hatarik, R., Hoover, D. E., Kilkenny, J. D., Kozioziemski, B. J., Kroll, J. J., McNaney, J. M., Nikroo, A., Sayre, D. B., Stadermann, M., Wild, C., Yoxall, B. E., Landen, O. L., Hsing, W. W., and Edwards, M. J. *Phys. Plasmas* **22**, 062703 (2015). 34

[53] Berzak Hopkins, L. F., Le Pape, S., Divol, L., Meezan, N. B., Mackinnon, A. J., Ho, D. D., Jones, O. S., Khan, S., Milovich, J. L., Ross, J. S., Amendt, P., Casey, D., Celliers, P. M., Pak, A., Peterson, J. L., Ralph, J., and Rygg, J. R. *Phys. Plasmas* **22**, 056318 (2015). 34

[54] Ross, J. S., Ho, D., Milovich, J., Döppner, T., McNaney, J., Macphee, A. G., Hamza, A., Biener, J., Robey, H. F., Dewald, E. L., Tommasini, R., Divol, L., Le Pape, S., Hopkins, L. B., Celliers, P. M., Landen, O., Meezan, N. B., and Mackinnon, A. J. *Phys. Rev. E - Stat. Nonlinear, Soft Matter Phys.* **91**, 021101 (2015). 34

[55] Divol, L., Pak, A., Berzak Hopkins, L. F., Le Pape, S., Meezan, N. B., Dewald, E. L., Ho, D. D., Khan, S. F., Mackinnon, A. J., Ross, J. S., Turnbull, D. P., Weber, C., Celliers, P. M., Millot, M., Benedetti, L. R., Field, J. E., Izumi, N., Kyrala, G. A., Ma, T., Nagel, S. R., Rygg, J. R., Edgell, D., Macphee, A. G., Goyon, C., Hohenberger, M., MacGowan, B. J., Michel, P., Strozzi, D., Cassata, W., Casey, D., Fittinghoff, D. N., Gharibyan, N., Hatarik, R., Sayre, D., Volegov, P., Yeamans, C., Bachmann, B., Döppner, T., Biener, J., Crippen, J., Choate, C., Huang, H., Kong, C., Nikroo, A., Rice, N. G., Stadermann, M., Bhandarkar, S. D., Haan, S., Kozioziemski, B., Hsing, W. W., Landen, O. L., Moody, J. D., Town, R. P., Callahan, D. A., Hurricane, O. A., and Edwards, M. J. *Phys. Plasmas* **24**, 056309 (2017). 34, 141, 150

[56] MacPhee, A. G., Casey, D. T., Clark, D. S., Felker, S., Field, J. E., Haan, S. W., Hammel, B. A., Kroll, J., Landen, O. L., Martinez, D. A., Michel, P., Milovich, J., Moore, A., Nikroo, A., Rice, N., Robey, H. F., Smalyuk, V. A., Stadermann, M., and Weber, C. R. *Phys. Rev. E* **95**, 031204 (2017). 34

[57] MacPhee, A. G., Smalyuk, V. A., Landen, O. L., Weber, C. R., Robey, H. F., Alfonso, E. L., Baker, K. L., Berzak Hopkins, L. F., Biener, J., Bunn, T., Casey, D. T., Clark, D. S., Crippen, J. W., Divol, L., Farrell, M., Felker, S., Field, J. E., Hsing, W. W., Kong, C., Le Pape, S., Martinez, D. A., Michel, P., Milovich, J.,

Moore, A., Nikroo, A., Pickworth, L., Rice, N., Stadermann, M., Yeamans, C., and Wild, C. *Phys. Plasmas* **25**(8), 082702 (2018). 34, 165, 171

[58] Callahan, D. A., Hurricane, O. A., Ralph, J. E., Thomas, C. A., Baker, K. L., Benedetti, L. R., Hopkins, L. F. B., Casey, D. T., Chapman, T., Czajka, C. E., Dewald, E. L., Divol, L., Döppner, T., Hinkel, D. E., Hohenberger, M., Khan, S. F., Kritcher, A. L., Landen, O. L., Lepape, S., Maclaren, S. A., Masse, L. P., and Meezan, N. B. *Phys. Plasmas* **25**, 056305 (2018). 35, 36

[59] Hinkel, D. E., Berzak Hopkins, L. F., Ma, T., Ralph, J. E., Albert, F., Benedetti, L. R., Celliers, P. M., Döppner, T., Goyon, C. S., Izumi, N., Jarrott, L. C., Khan, S. F., Kline, J. L., Kritcher, A. L., Kyrala, G. A., Nagel, S. R., Pak, A. E., Patel, P., Rosen, M. D., Rygg, J. R., Schneider, M. B., Turnbull, D. P., Yeamans, C. B., Callahan, D. A., and Hurricane, O. A. *Phys. Rev. Lett.* **117**(22), 225002 (2016). 35

[60] Clark, D. S., Kritcher, A. L., Yi, S. A., Zylstra, A. B., Haan, S. W., and Weber, C. R. *Phys. Plasmas* **25**, 032703 (2018). 35

[61] Kritcher, A. L., Clark, D. S., Haan, S., Yi, S. A., Zylstra, A. B., Callahan, D. A., Hinkel, D. E., Berzak Hopkins, L. F., Hurricane, O. A., Landen, O. L., Maclaren, S. A., Meezan, N. B., Patel, P. K., Ralph, J., Thomas, C. A., Town, R., and Edwards, J. *Phys. Plasmas* **25**, 056309 (2018). 35

[62] Döppner, T., Callahan, D. A., Hurricane, O. A., Hinkel, D. E., Ma, T., Park, H. S., Berzak Hopkins, L. F., Casey, D. T., Celliers, P., Dewald, E. L., Dittrich, T. R., Haan, S. W., Kritcher, A. L., MacPhee, A., Le Pape, S., Pak, A., Patel, P. K., Springer, P. T., Salmonson, J. D., Tommasini, R., Benedetti, L. R., Bond, E., Bradley, D. K., Caggiano, J., Church, J., Dixit, S., Edgell, D., Edwards, M. J., Fittinghoff, D. N., Frenje, J., Gatu Johnson, M., Grim, G., Hatarik, R., Havre, M., Herrmann, H., Izumi, N., Khan, S. F., Kline, J. L., Knauer, J., Kyrala, G. A., Landen, O. L., Merrill, F. E., Moody, J., Moore, A. S., Nikroo, A., Ralph, J. E., Remington, B. A., Robey, H. F., Sayre, D., Schneider, M., Streckert, H., Town, R., Turnbull, D., Volegov, P. L., Wan, A., Widmann, K., Wilde, C. H., and Yeamans, C. *Phys. Rev. Lett.* **115**, 055001 (2015). 35

[63] Leidinger, J.-P., Callahan, D. A., Berzak-Hopkins, L. F., Ralph, J. E., Amendt, P., Hinkel, D. E., Michel, P., Moody, J. D., Ross, J. S., Rygg, J. R., Celliers, P., Clouët, J.-F., Dewald, E. L., Kaiser, P., Khan, S., Kritcher, A. L., Liberatore, S., Marion, D., Masson-Laborde, P.-E., Milovich, J. L., Morice, O., Pak, A. E., Poujade, O., Strozzi, D., and Hurricane, O. A. *J. Phys. Conf. Ser.* **717**, 012035 (2016). 35

[64] Farmer, W. A., Tabak, M., Hammer, J. H., Amendt, P. A., and Hinkel, D. E. *Phys. Plasmas* **26**(3), 032701 (2019). 35

[65] Clark, D. S., Weber, C. R., Smalyuk, V. A., Robey, H. F., Kritcher, A. L., Milovich, J. L., and Salmonson, J. D. *Phys. Plasmas* **23**, 072707 (2016). 35

[66] Baker, K. L., Robey, H. F., Milovich, J. L., Jones, O. S., Smalyuk, V. A., Casey, D. T., MacPhee, A. G., Pak, A., Celliers, P. M., Clark, D. S., Landen, O. L., Peterson, J. L., Berzak-Hopkins, L. F., Weber, C. R., Haan, S. W., Döppner, T. D., Dixit, S., Giraldez, E., Hamza, A. V., Jancaitis, K. S., Kroll, J. J., Lafortune, K. N., MacGowan, B. J., Moody, J. D., Nikroo, A., and Widmayer, C. C. *Phys. Plasmas* **22**, 052702 (2015). 35

[67] Milovich, J. L., Robey, H. F., Clark, D. S., Baker, K. L., Casey, D. T., Cerjan, C., Field, J., MacPhee, A. G., Pak, A., Patel, P. K., Peterson, J. L., Smalyuk, V. A., and Weber, C. R. *Phys. Plasmas* **22**, 122702 (2015). 35

[68] Peterson, J. L., Berzak Hopkins, L. F., Jones, O. S., and Clark, D. S. *Phys. Rev. E - Stat. Nonlinear, Soft Matter Phys.* **91**, 031101 (2015). 35

[69] Smalyuk, V. A., Robey, H. F., Döppner, T. D., Jones, O. S., Milovich, J. L., Bachmann, B., Baker, K. L., Berzak Hopkins, L. F., Bond, E., Callahan, D. A., Casey, D. T., Celliers, P. M., Cerjan, C., Clark, D. S., Dixit, S. N., Edwards, M. J., Giraldez, E., Haan, S. W., Hamza, A. V., Hohenberger, M., Hoover, D., Hurricane, O. A., Jancaitis, K. S., Kroll, J. J., Lafortune, K. N., Landen, O. L., MacGowan, B. J., MacPhee, A. G., Nikroo, A., Pak, A., Patel, P. K., Peterson, J. L., Weber, C. R., Widmayer, C. C., and Yeamans, C. *Phys. Plasmas* **22**, 080703 (2015). 35

[70] MacPhee, A. G., Peterson, J. L., Casey, D. T., Clark, D. S., Haan, S. W., Jones, O. S., Landen, O. L., Milovich, J. L., Robey, H. F., and Smalyuk, V. A. *Phys. Plasmas* **22**, 080702 (2015). 35

[71] Smalyuk, V. A., Robey, H. F., Döppner, T., Casey, D. T., Clark, D. S., Jones, O. S., Milovich, J. L., Peterson, J. L., Bachmann, B., Baker, K. L., Benedetti, L. R., Berzak Hopkins, L. F., Bionta, R., Bond, E., Bradley, D. K., Callahan, D. A., Celliers, P. M., Cerjan, C., Chen, K. C., Goyon, C., Grim, G., Dixit, S. N., Eckart, M. J., Edwards, M. J., Farrell, M., Fittinghoff, D. N., Frenje, J. A., Gatu-Johnson, M., Gharibyan, N., Haan, S. W., Hamza, A. V., Hartouni, E., Hatarik, R., Havre, M., Hohenberger, M., Hoover, D., Hurricane, O. A., Izumi, N., Jancaitis, K. S., Khan, S. F., Knauer, J. P., Kroll, J. J., Kyrala, G., LaFortune, K. N., Landen, O. L., Ma, T., MacGowan, B. J., MacPhee, A. G., Mauldin, M., Merrill, F. E., Moore, A. S., Nagel, S., Nikroo, A., Pak, A., Patel, P. K., Ralph, J. E., Sayre, D. B., Shaughnessy, D., Spears, B. K., Tommasini, R., Turnbull, D. P., Velikovich, A. L., Volegov, P. L., Weber, C. R., Widmayer, C. C., and Yeamans, C. *Phys. Plasmas* **23**, 102703 (2016). 35

[72] Robey, H. F., Smalyuk, V. A., Milovich, J. L., Döppner, T. D., Casey, D. T., Baker, K. L., Peterson, J. L., Bachmann, B., Berzak Hopkins, L. F., Bond, E.,

Caggiano, J. A., Callahan, D. A., Celliers, P. M., Cerjan, C., Clark, D. S., Dixit, S. N., Edwards, M. J., Gharibyan, N., Haan, S. W., Hammel, B. A., Hamza, A. V., Hatarik, R., Hurricane, O. A., Jancaitis, K. S., Jones, O. S., Kerbel, G. D., Kroll, J. J., Lafortune, K. N., Landen, O. L., Ma, T., Marinak, M. M., Macgowan, B. J., Macphee, A. G., Pak, A., Patel, M., Patel, P. K., Perkins, L. J., Sayre, D. B., Sepke, S. M., Spears, B. K., Tommasini, R., Weber, C. R., Widmayer, C. C., Yeamans, C., Giraldez, E., Hoover, D., Nikroo, A., Hohenberger, M., and Gatu Johnson, M. *Phys. Plasmas* **23**, 056303 (2016). 35, 173

[73] Xu, H. W., Alford, C. S., Cooley, J. C., Dixon, L. A., Hackenberg, R. E., Letts, S. A., Moreno, K. A., Nikroo, A., Wall, J. R., and Youngblood, K. P. *Fusion Sci. Technol.* **51**(4), 547–552 (2007). 35

[74] Kline, J. L., Yi, S. A., Simakov, A. N., Olson, R. E., Wilson, D. C., Kyrala, G. A., Perry, T. S., Batha, S. H., Zylstra, A. B., Dewald, E. L., Tommasini, R., Ralph, J. E., Strozzi, D. J., MacPhee, A. G., Callahan, D. A., Hinkel, D. E., Hurricane, O. A., Milovich, J. L., Rygg, J. R., Khan, S. F., Haan, S. W., Celliers, P. M., Clark, D. S., Hammel, B. A., Kozioziemski, B., Schneider, M. B., Marinak, M. M., Rinderknecht, H. G., Robey, H. F., Salmonson, J. D., Patel, P. K., Ma, T., Edwards, M. J., Stadermann, M., Baxamusa, S., Alford, C., Wang, M., Nikroo, A., Rice, N., Hoover, D., Youngblood, K. P., Xu, H., Huang, H., and Sio, H. *Phys. Plasmas* **23**, 056310 (2016). 35

[75] Simakov, A. N., Wilson, D. C., Yi, S. A., Loomis, E. N., Kline, J. L., Kyrala, G. A., Zylstra, A. B., Dewald, E. L., Tommasini, R., Ralph, J. E., Strozzi, D. J., Macphee, A. G., Milovich, J. L., Rygg, J. R., Khan, S. F., Ma, T., Jarrott, L. C., Haan, S. W., Celliers, P. M., Marinak, M. M., Rinderknecht, H. G., Robey, H. F., Salmonson, J. D., Stadermann, M., Baxamusa, S., Alford, C., Wang, Y., Nikroo, A., Rice, N., Kong, C., Jaquez, J., Mauldin, M., Youngblood, K. P., Xu, H., Huang, H., and Sio, H. *Phys. Plasmas* **24**, 052704 (2017). 35

[76] Zylstra, A. B., Yi, S. A., Maclaren, S., Kline, J., Kyrala, G. A., Ralph, J. E., Bae, J., Batha, S., Callahan, D. A., Flippo, K., Huang, H., Hurricane, O. A., Khan, S. F., Kabadi, N., Kong, C., Kot, L. B., Lahmann, B., Loomis, E. N., Masse, L. P., Millot, M., Moore, A., Nikroo, A., Perry, T. S., Rice, N., Salmonson, J., Shah, R., Sio, H., Stadermann, M., Strozzi, D. J., Tipton, R., Xu, H., Zylstra, A. B., Yi, S. A., Maclaren, S., Kline, J., Kyrala, G., Ralph, J. E., and Bae, J. *Phys. Plasmas* **25**, 102704 (2018). 35

[77] Weber, C. R., Casey, D. T., Clark, D. S., Hammel, B. A., MacPhee, A., Milovich, J., Martinez, D., Robey, H. F., Smalyuk, V. A., Stadermann, M., Amendt, P., Bhandarkar, S., Chang, B., Choate, C., Crippen, J., Felker, S. J., Field, J. E., Haan, S. W., Johnson, S., Kroll, J. J., Landen, O. L., Marinak, M., Mcinnis, M., Nikroo, A., Rice, N., and Sepke, S. M. *Phys. Plasmas* **24**, 056302 (2017). 36

[78] Smalyuk, V. A., Robey, H. F., Alday, C. L., Amendt, P., Aracne-Ruddle, C., Bigelow, J. R., Bunn, T., Casey, D. T., Chen, K.-C., Clark, D. S., Cortez, J. P., Crippen, J., Diaz, S., Farrell, M., Felker, S., Field, J. E., Jaquez, J., Johnson, S., Haan, S. W., Hammel, B. A., Hamza, A. V., Havre, M. O., Heinbockel, C., Hsing, W. W., Kangas, K., Kroll, J. J., Kucheyev, S. O., Landen, O. L., Lepro-Chavez, X., MacPhee, A. G., Martinez, D. A., Milovich, J., Nikroo, A., Pickworth, L. A., Rice, N., Stadermann, M., Steich, D., and Weber, C. R. *Phys. Plasmas* **25**(7), 072705 (2018). 36

[79] Robey, H. F., Berzak Hopkins, L., Milovich, J. L., and Meezan, N. B. *Phys. Plasmas* **25**, 012711 (2018). 36

[80] Lobatchev, V. and Betti, R. *Phys. Rev. Lett.* **85**(21), 4522–5 (2000). 37, 53

[81] Atzeni, S. *Comput. Phys. Commun.* **43**(1), 107–124 (1986). 37

[82] Fan, Z., Zhu, S., Pei, W., Ye, W., Li, M., Xu, X., Wu, J., Dai, Z., and Wang, L. *EPL* **99**, 65003 (2012). 37

[83] Levedahl, W. K. and Lindl, J. D. *Nucl. Fusion* **37**(2), 165–173 (1997). 37, 55

[84] Kishony, R. and Shvarts, D. *Phys. Plasmas* **8**(11), 4925–4936 (2001). 37, 40, 54, 56

[85] Springer, P. T., Hurricane, O. A., Hammer, J. H., Betti, R., Callahan, D. A., Campbell, E. M., Casey, D. T., Cerjan, C. J., Cao, D., Dewald, E. L., Divol, L., Doeppner, T., Edwards, M. J., Field, J. E., Forrest, C. J., Frenje, J. A., Gaffney, J. A., Gatu-Johnson, M., Glebov, V. Y., Goncharov, V. N., Grim, G. P., Hartouni, E. P., Hatarik, R., Hinkel, D. E., Berzak Hopkins, L. F., Igumenshchev, I. V., Knapp, P. F., Knauer, J. P., Kritcher, A. L., Landen, O. L., Pak, A. E., Le Pape, S., Ma, T., MacPhee, A. G., Munro, D. H., Nora, R. C., Patel, P. K., Peterson, L., Radha, P. B., Regan, S. P., Rinderknecht, H. G., Sangster, T. C., Spears, B. K., and Stoeckl, C. *Nucl. Fusion* **59**, 032009 (2019). 39, 56, 118, 168

[86] Cheng, B., Kwan, T. J., Yi, S. A., Landen, O. L., Wang, Y. M., Cerjan, C. J., Batha, S. H., and Wysocki, F. J. *Phys. Rev. E* **98**, 023203 (2018). 39

[87] Cheng, B., Kwan, T. J., Wang, Y. M., and Batha, S. H. *Phys. Plasmas* **21**, 102707 (2014). 39, 56, 118, 168

[88] Cheng, B., Kwan, T. J., Wang, Y. M., Merrill, F. E., Cerjan, C. J., and Batha, S. H. *Phys. Plasmas* **22**, 082704 (2015). 39

[89] Taylor, S. and Chittenden, J. P. *Phys. Plasmas* **21**(6), 062701 (2014). 39, 40, 50, 54, 106

[90] Bose, A., Betti, R., Shvarts, D., and Woo, K. M. *Phys. Plasmas* **24**, 102704 (2017). 40

[91] Taylor, S. *A Computational Study of the Stagnation Phase in Inertial Confinement Fusion : Hotspot Energetics, Diagnostics, and Burn Propagation.* PhD thesis, Imperial College London, (2013). 40, 51, 87, 106, 121, 122, 139, 165

[92] Spitzer, L. *Physics of Fully Ionized Gases.* Wiley Interscience, 2nd ed. edition, (1962). 42, 43, 50

[93] Cohen, R. S., Spitzer, L., and Routly, P. M. *Phys. Rev.* **80**(2), 230–238 (1950). 43

[94] Chandrasekhar, S. (1943). 43

[95] Li, C. K. and Petrasso, R. D. *Phys. Rev. Lett.* **70**(20), 3063–3066 (1993). 44

[96] Li, C.-K. and Petrasso, R. D. *Phys. Rev. Lett.* **70**(20), 3059–3062 (1993). 44, 79

[97] Li, C.-K. and Petrasso, R. D. *Phys. Rev. Lett.* **114**(19), 199901 (2015). 44, 48

[98] Maynard, G. and Deutsch, C. *J. Phys.* **46**(7), 1113–1122 (1985). 44

[99] Gouedard, C. and Deutsch, C. *J. Math. Phys.* **19**(1), 32 (1978). 44

[100] Zimmerman, G. B. Technical Report UCRL-JC-126S45, Lawrence Livermore National Laboratory, (1997). 44

[101] Gericke, D. O., Murillo, M. S., and Schlanges, M. *Phys. Rev. E - Stat. Nonlinear, Soft Matter Phys.* **65**, 036418 (2002). 45

[102] Brown, L. S., Preston, D. L., and Singleton, R. L. *Phys. Rep.* **410**(4), 237–333 (2005). 45, 46, 79

[103] Singleton, R. L. *Phys. Plasmas* **15**, 056302 (2008). 46, 79

[104] Gericke, D. O. and Schlanges, M. *Phys. Rev. E - Stat. Physics, Plasmas, Fluids, Relat. Interdiscip. Top.* **60**(1), 904–910 (1999). 46

[105] Gericke, D. O. *Laser Part. Beams* **20**(04), 471–474 (2002).

[106] Gericke, D. and Schlanges, M. *Phys. Rev. E* **67**(3), 037401 (2003). 46, 48

[107] Hayes, A. C. Technical Report LA-UR-13-22639, Los Alamos National Laboratory, (2013). 47, 48

[108] Hayes, A. C., Jungman, G., Schulz, A. E., Boswell, M., Fowler, M. M., Grim, G., Klein, A., Rundberg, R. S., Wilhelmy, J. B., and Wilson, D. *Phys. Plasmas* **22**, 082703 (2015). 47, 174

[109] Frenje, J. A., Grabowski, P. E., Li, C. K., Séguin, F. H., Zylstra, A. B., Gatu Johnson, M., Petrasso, R. D., Glebov, V. Y., and Sangster, T. C. *Phys. Rev. Lett.* **115**, 205001 (2015). 48

[110] Cayzac, W., Frank, A., Ortner, A., Bagnoud, V., Basko, M. M., Bedacht, S., Bläser, C., Blažević, A., Busold, S., Deppert, O., Ding, J., Ehret, M., Fiala, P., Frydrych, S., Gericke, D. O., Hallo, L., Helfrich, J., Jahn, D., Kjartansson, E., Knetsch, A., Kraus, D., Malka, G., Neumann, N. W., Pépitone, K., Pepler, D., Sander, S., Schaumann, G., Schlegel, T., Schroeter, N., Schumacher, D., Seibert, M., Tauschwitz, A., Vorberger, J., Wagner, F., Weih, S., Zobus, Y., and Roth, M. *Nat. Commun.* **8**(June), 15693 (2017). 48

[111] Betti, R., Umansky, M., Lobatchev, V., Goncharov, V. N., and McCrory, R. L. *Phys. Plasmas* **8**(12), 5257 (2001). 50, 53, 102

[112] Zel'dovich, Y. and Raizer, Y. *Physics of Shock Waves and High-Temperature Hydrodynamic Phenomena.* Dover Books on Physics. Dover Publications, (2002). 50, 51

[113] Gauthier, P., Chaland, F., and Masse, L. *Phys. Rev. E - Stat. Nonlinear, Soft Matter Phys.* **70**, 055401 (2004). 51, 58, 84

[114] Epperlein, E. M. and Haines, M. G. *Phys. Fluids* **29**(4), 1029 (1986). 51, 62

[115] Chandrasekhar, S. *Hydrodynamic and Hydromagnetic Stability.* Dover Books on Physics Series. Dover Publications, (1961). 53

[116] Bodner, S. E. *Phys. Rev. Lett.* **33**(13), 761–764 (1974). 53

[117] Takabe, H., Mima, K., Montierth, L., and Morse, R. L. *Phys. Fluids* **28**(12), 3676 (1985). 53

[118] Goncharov, V. N., Betti, R., McCrory, R. L., Sorotokin, P., and Verdon, C. P. *Phys. Plasmas* **3**(4), 1402 (1996).

[119] Betti, R., Goncharov, V. N., McCrory, R. L., and Verdon, C. P. *Phys. Plasmas* **5**(5), 1446 (1998). 53

[120] Lindl, J. *Inertial Confinement Fusion: The Quest for Ignition and Energy Gain Using Indirect Drive.* AIP-Press Series. AIP Press, (1998). 53

[121] Lawson, J. D. *Proc. Phys. Soc.* **70**, 6–10 (1957). 55

[122] Basko, M. *Nucl. Fusion* **35**, 87 January (1995). 55

[123] Herrmann, M. C., Tabak, M., and Lindl, J. D. *Nucl. Fusion* **41**(1), 99–111 (2001). 55

[124] Betti, R., Chang, P. Y., Spears, B. K., Anderson, K. S., Edwards, J., Fatenejad, M., Lindl, J. D., McCrory, R. L., Nora, R., and Shvarts, D. *Phys. Plasmas* **17**, 058102 (2010). 55

[125] Cheng, B., Kwan, T. J., Wang, Y. M., and Batha, S. H. *Phys. Rev. E - Stat. Nonlinear, Soft Matter Phys.* **88**, 041101 (2013). 56, 118

[126] Krokhin, O. N. and Rozanov, V. B. *Sov. J. Quantum Electron.* **2**(4), 393–394 (1973). 57, 100

[127] Atzeni, S. and Caruso, A. *Nuovo Cim. B* **80**(1), 71–103 (1984). 57, 100

[128] Nora, R., Betti, R., Anderson, K. S., Shvydky, A., Bose, A., Woo, K. M., Christopherson, A. R., Marozas, J. A., Collins, T. J., Radha, P. B., Hu, S. X., Epstein, R., Marshall, F. J., McCrory, R. L., Sangster, T. C., and Meyerhofer, D. D. *Phys. Plasmas* **21**, 056316 (2014). 58, 59, 142, 146, 162

[129] Zhou, C. D. and Betti, R. *Phys. Plasmas* **14**, 072703 (2007). 59, 162

[130] McGlinchey, K. *Radiation-hydrodynamic simulations of the impact of instabilities and asymmetries on inertial confinement fusion.* PhD thesis, Imperial College London, (2017). 60, 62, 63, 88, 116, 120, 121, 139, 143, 151

[131] Walsh, C. A. *Extended Magneto-hydrodynamic Effects in Indirect-Drive Inertial Confinement Fusion Experiments.* PhD thesis, Imperial College London, (2018). 60, 62, 172, 173

[132] Van Leer, B. *J. Comput. Phys.* **32**, 101–136 (1979). 61

[133] Youngs, D. L. In *Numer. Methods Fluid Dyn.,* Morton, K. W. and Baines, M. J., editors. Academic Press, New York (1982). 61

[134] Noh, W. F. and Woodward, P. In *Proc. Fifth Int. Conf. Numer. Methods Fluid Dyn.*, 330–340. Lecture Notes in Physics, Vol. 59 Springer, Berlin, Heidelburg (1976). 61

[135] Faik, S., Basko, M. M., Tauschwitz, A., Iosilevskiy, I., and Maruhn, J. A. *High Energy Density Phys.* **8**(4), 349–359 (2012). 61

[136] Kemp, A. J. and Meyer-ter Vehn, J. *Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip.* **415**(3), 674–676 (1998).

[137] More, R. M., Warren, K. H., Young, D. A., and Zimmerman, G. B. *Phys. Fluids* **31**(10), 3059–3078 (1988). 61

[138] Braginskii, S. I. *Rev. Plasma Phys.* **1**, 205 (1965). 62

[139] Meyer, C. D., Balsara, D. S., and Aslam, T. D. *J. Comput. Phys.* **257**(PA), 594–626 (2014). 62

[140] Olson, G., Auer, L., and Hall, M. *J. Quant. Spectrosc. Radiat. Transf.* **64**, 619–634 (2000). 62, 179

[141] Mihalas, D. and Weaver, R. *J. Quant. Spectrosc. Radiat. Transf.* **28**(3), 213–222 (1982). 62

[142] Jennings, C. A. *Radiation Transport Effects in Wire Array Z-Pinches and Magneto-Hydrodynamic Modelling Techniques.* PhD thesis, Imperial College London, (2005). 62

[143] Niasse, N.-P. L. *Development of a Pseudo Non-LTE model for Z-pinch simulations.* PhD thesis, Imperial College London, (2011). 63

[144] Faussurier, G., Blancard, C., and Renaudin, P. *High Energy Density Phys.* **4**(3-4), 114–123 (2008). 63

[145] Kramida, A., Ralchenko, Y., Reader, J., and NIST ASD Team. *NIST Atomic Spectra Database* , National Institution of Standards and Technology (2018). 63

[146] Busquet, M., Colombant, D., Klapisch, M., Fyfe, D., and Gardner, J. *High Energy Density Phys.* **5**(4), 270–275 (2009). 63

[147] Corman, E. G., Loewe, W. E., Cooper, G. E., and Winslow, A. M. *Nucl. Fusion* **15**, 377–86 (1975). 64

[148] Mehlhorn, T. A. and Duderstadt, J. J. *J. Comput. Phys.* **38**(1), 86–106 (1980). 64

[149] Morel, J. E. *Nucl. Sci. Eng.* **79**(4), 340–356 (1981).

[150] Honrubia, J. J. and Aragones, J. M. *Nucl. Sci. Eng.* **93**, 386–402 (1986). 64

[151] Lathrop, K. D. *Nucl. Sci. Eng.* **32**(3), 357–369 June (1968). 64

[152] Takizuka, T. and Abe, H. *J. Comput. Phys.* **25**(3), 205–219 (1977). 64

[153] Nanbu, K. *Phys. Rev. E* **55**(4), 4642–4652 (1997). 64

[154] Cadjan, M. G. and Ivanov, M. F. *J. Plasma Phys.* **61**(1), 89–106 (1999). 64

[155] Jones, M. E., Lemons, D. S., Mason, R. J., Thomas, V. A., and Winske, D. *J. Comput. Phys.* **123**(1), 169–181 (1996). 64

[156] Peigney, B. E., Larroche, O., and Tikhonchuk, V. *Phys. Plasmas* **21**, 122709 (2014). 65

[157] Peigney, B. E., Larroche, O., and Tikhonchuk, V. *J. Comput. Phys.* **278**, 416–444 (2014). 65

[158] Sherlock, M. *J. Comput. Phys.* **227**(4), 2286–2292 February (2008). 65, 66, 80

[159] Birdsall, C. K. and Langdon, A. B. *Plasma Physics via Computer Simulation.* CRC Press, 1st. ed. edition, (1985). 66

[160] Huba, J. D. Technical Report NRL/PU/6790–18-640, Naval Research Laboratory, Washington, D.C., (2013). 66

[161] Boris, J. P. *Proc. Fourth Conf. Numer. Simul. Plasmas* , 3–67 (1970). 67

[162] Bosch, H.-S. and Hale, G. *Nucl. Fusion* **33**(12), 1919–1919 (1992). 68

[163] Appelbe, B. and Chittenden, J. *Plasma Phys. Control. Fusion* **53**(53), 45002–45002 (2011). 68

[164] Clayton, D. D. *Principles of Stellar Evolution and Nucleosynthesis*. The University of Chicago Press, 2nd. ed. edition, (1983). 68

[165] Vickers, S. W. *Particle in Cell and Hybrid Simulations of the Z Double-Post-Hole Convolute Cathode Plasma Evolution and Dynamics*. PhD thesis, Imperial College London, (2012). 68

[166] Knuth, D. *The Art of Computer Programming*, volume 2. Addison Wesley Professional, 3rd. ed. edition, (1998). 69, 70

[167] Press, W., Teukolsky, S., Vetterling, W., and Flannery, B. *Numerical Recipes 3rd Edition: The Art of Scientific Computing*. Cambridge University Press, (2007). 69

[168] Katzgraber, H. G. *arXiv Prepr. arXiv1005.4117* , 1–20 (2010). 70

[169] L'Ecuyer, P. and Simard, R. *ACM Trans. Math. Softw.* **33**(4), 22–es (2007). 71

[170] De Matteis, A. and Pagnutti, S. *Parallel Comput.* **13**(2), 193–198 (1990). 71

[171] Bauke, H. and Mertens, S. *Phys. Rev. E - Stat. Nonlinear, Soft Matter Phys.* **75**(6), 1–14 (2007). 72

[172] Matsumoto, M. and Nishimura, T. *ACM Trans. Model. Comput. Simul.* **8**(1), 3–30 (1998). 73

[173] Marsaglia, G. *J. Stat. Softw.* **8**(14), 1–6 (2003). 74

[174] Panneton, F. and L'ecuyer, P. *ACM Trans. Model. Comput. Simul.* **15**(4), 346–361 (2005). 74

[175] Vigna, S. *ACM Trans. Math. Softw.* **42**(4), 30 (2016). 74

[176] Vigna, S. *J. Comput. Appl. Math.* **315**, 175–181 (2017). 74

[177] Thomas, D. B., Luk, W., Leong, P. H. W., and Villasenor, J. D. *ACM Comput. Surv.* **39**(4), 11 (2007). 74

[178] Marsaglia, G. and Tsang, W. W. *J. Stat. Softw.* **5**(8), 1–7 (2000). 75

[179] Atzeni, S. *Plasma Phys. Control. Fusion* **29**(11), 1535–1604 (1987). 87

[180] Betti, R., Christopherson, A. R., Spears, B. K., Nora, R., Bose, A., Howard, J., Woo, K. M., Edwards, M. J., and Sanz, J. *Phys. Rev. Lett.* **114**(25), 255003 (2015). 90

[181] Rybicki, G. and Lightman, A. *Radiative Processes in Astrophysics.* A Wiley-Interscience publication. Wiley, (1979). 93

[182] Chittenden, J. P., Appelbe, B. D., Manke, F., McGlinchey, K., and Niasse, N. P. L. *Phys. Plasmas* **23**(5), 052708 (2016). 106

[183] Grim, G. P., Guler, N., Merrill, F. E., Morgan, G. L., Danly, C. R., Volegov, P. L., Wilde, C. H., Wilson, D. C., Clark, D. S., Hinkel, D. E., Jones, O. S., Raman, K. S., Izumi, N., Fittinghoff, D. N., Drury, O. B., Alger, E. T., Arnold, P. A., Ashabranner, R. C., Atherton, L. J., Barrios, M. A., Batha, S., Bell, P. M., Benedetti, L. R., Berger, R. L., Bernstein, L. A., Berzins, L. V., Betti, R., Bhandarkar, S. D., Bionta, R. M., Bleuel, D. L., Boehly, T. R., Bond, E. J., Bowers, M. W., Bradley, D. K., Brunton, G. K., Buckles, R. A., Burkhart, S. C., Burr, R. F., Caggiano, J. A., Callahan, D. A., Casey, D. T., Castro, C., Celliers, P. M., Cerjan, C. J., Chandler, G. A., Choate, C., Cohen, S. J., Collins, G. W., Cooper, G. W., Cox, J. R., Cradick, J. R., Datte, P. S., Dewald, E. L., Di Nicola, P., Di Nicola, J. M., Divol, L., Dixit, S. N., Dylla-Spears, R., Dzenitis, E. G., Eckart, M. J., Eder, D. C., Edgell, D. H., Edwards, M. J., Eggert, J. H., Ehrlich, R. B., Erbert, G. V., Fair, J., Farley, D. R., Felker, B., Fortner, R. J., Frenje, J. A., Frieders, G., Friedrich, S., Gatu-Johnson, M., Gibson, C. R., Giraldez, E., Glebov, V. Y., Glenn, S. M., Glenzer, S. H., Gururangan, G., Haan, S. W., Hahn, K. D., Hammel, B. A., Hamza, A. V., Hartouni, E. P., Hatarik, R., Hatchett, S. P., Haynam, C., Hermann, M. R., Herrmann, H. W., Hicks, D. G., Holder, J. P., Holunga, D. M., Horner, J. B., Hsing, W. W., Huang, H., Jackson, M. C., Jancaitis, K. S., Kalantar, D. H., Kauffman, R. L., Kauffman, M. I., Khan, S. F., Kilkenny, J. D., Kimbrough, J. R., Kirkwood, R., Kline, J. L., Knauer, J. P., Knittel, K. M., Koch, J. A., Kohut, T. R., Kozioziemski, B. J., Krauter, K., Krauter, G. W., Kritcher, A. L., Kroll, J., Kyrala, G. A., La Fortune, K. N., Lacaille, G., Lagin, L. J., Land, T. A., Landen, O. L., Larson, D. W., Latray, D. A., Leeper, R. J., Lewis, T. L., Lepape, S., Lindl, J. D., Lowe-Webb, R. R., Ma, T., Macgowan, B. J., Mackinnon, A. J., Macphee, A. G., Malone, R. M., Malsbury, T. N., Mapoles, E., Marshall, C. D., Mathisen, D. G., McKenty, P., McNaney, J. M., Meezan, N. B., Michel, P., Milovich, J. L., Moody, J. D., Moore, A. S., Moran, M. J., Moreno, K., Moses, E. I., Munro, D. H., Nathan, B. R., Nelson, A. J., Nikroo, A., Olson, R. E., Orth, C., Pak, A. E., Palma, E. S., Parham, T. G., Patel, P. K., Patterson, R. W., Petrasso, R. D., Prasad, R., Ralph, J. E., Regan, S. P., Rinderknecht, H., Robey, H. F., Ross, G. F., Ruiz, C. L., Séguin, F. H., Salmonson, J. D., Sangster, T. C., Sater, J. D., Saunders, R. L., Schneider, M. B., Schneider, D. H., Shaw, M. J., Simanovskaia, N., Spears, B. K., Springer,

P. T., Stoeckl, C., Stoeffl, W., Suter, L. J., Thomas, C. A., Tommasini, R., Town, R. P., Traille, A. J., Wonterghem, B. V., Wallace, R. J., Weaver, S., Weber, S. V., Wegner, P. J., Whitman, P. K., Widmann, K., Widmayer, C. C., Wood, R. D., Young, B. K., Zacharias, R. A., and Zylstra, A. *Phys. Plasmas* **20**, 056320 (2013). 106

[184] Christopherson, A. R., Betti, R., Bose, A., Howard, J., Woo, K. M., Campbell, E. M., Sanz, J., and Spears, B. K. *Phys. Plasmas* **25**, 012703 (2018). 106

[185] Christopherson, A. R., Betti, R., Howard, J., Woo, K. M., Bose, A., Campbell, E. M., and Gopalaswamy, V. *Phys. Plasmas* **25**, 072704 (2018). 106

[186] Layzer, D. *Astrophys. J.* **122**, 1 (1955). 121, 139, 165

[187] Hammel, B. A., Weber, C. R., Stadermann, M., Alday, C. L., Aracne-Ruddle, C., Bigelow, J. R., Clark, D. S., Cortez, J. P., Diaz, S., Doeppner, T., Felker, S. J., Field, J. E., Haan, S. W., Havre, M. O., Heinbockel, C., Hinkel, D. E., Hsing, W. W., Johnson, S. A., Nikroo, A., Pickworth, L. A., Ralph, J. E., Robey, H. F., and Smalyuk, V. A. *Phys. Plasmas* **25**, 082714 (2018). 133

[188] Casey, D. T., Thomas, C. A., Baker, K. L., Spears, B. K., Hohenberger, M., Khan, S. F., Nora, R. C., Weber, C. R., Woods, D. T., Hurricane, O. A., Callahan, D. A., Berger, R. L., Milovich, J. L., Patel, P. K., Ma, T., Pak, A., Benedetti, L. R., Millot, M., Jarrott, C., Landen, O. L., Bionta, R. M., MacGowan, B. J., Strozzi, D. J., Stadermann, M., Biener, J., Nikroo, A., Goyon, C. S., Izumi, N., Nagel, S. R., Bachmann, B., Volegov, P. L., Fittinghoff, D. N., Grim, G. P., Yeamans, C. B., Gatu Johnson, M., Frenje, J. A., Rice, N., Kong, C., Crippen, J., Jaquez, J., Kangas, K., and Wild, C. *Phys. Plasmas* **25**, 056308 (2018). 141

[189] Hurricane, O., Callahan, D. A., Casey, D. T., Celliers, P. M., Cerjan, C., Dewald, E. L., Dittrich, T. R., Döppner, T., Hinkel, D. E., Berzak Hopkins, L. F., Kline, J. L., Le Pape, S., Ma, T., Macphee, A. G., Milovich, J. L., Pak, A., Park, H.-S., Patel, P. K., Remington, B. A., Salmonson, J., Springer, P. T., and Tommasini, R. *Nature* **506**(7488), 343–8 (2014). 141

[190] Hurricane, O. A., Kritcher, A., Callahan, D. A., Landen, O., Patel, P. K., Springer, P. T., Casey, D. T., Dewald, E. L., Dittrich, T. R., Döppner, T., Hinkel, D. E., Berzak Hopkins, L. F., Kline, J., Le Pape, S., Ma, T., MacPhee, A. G., Moore, A., Pak, A., Park, H. S., Ralph, J., Salmonson, J. D., and Widmann, K. *Phys. Plasmas* **24**, 092706 (2017). 151

[191] Haan, S. W., Kritcher, A. L., Clark, D. S., Yi, S. A., Zylstra, A. B., Ralph, J. E., and Weber, C. R. Technical Report LLNL-TR-741418, Lawrence Livermore National Laboratory, (2017). 152

[192] Desai, R. *Personal Communication* , Imperial College London (2019). 172

[193] Walsh, C. A., Mcglinchey, K., Tong, J. K., Appelbe, B. D., Crilly, A. J., Zhang, M. F., and Chittenden, J. P. *Phys. Plasmas* **26**, 022701 (2019). 173

[194] Döppner, T., Thomas, C. A., Divol, L., Dewald, E. L., Celliers, P. M., Bradley, D. K., Callahan, D. A., Dixit, S. N., Harte, J. A., Glenn, S. M., Haan, S. W., Izumi, N., Kyrala, G. A., LaCaille, G., Kline, J. K., Kruer, W. L., Ma, T., MacKinnon, A. J., McNaney, J. M., Meezan, N. B., Robey, H. F., Salmonson, J. D., Suter, L. J., Zimmerman, G. B., Edwards, M. J., MacGowan, B. J., Kilkenny, J. D., Lindl, J. D., Van Wonterghem, B. M., Atherton, L. J., Moses, E. I., Glenzer, S. H., and Landen, O. L. *Phys. Rev. Lett.* **108**, 135006 (2012). 173

[195] Crilly, A. J., Appelbe, B. D., McGlinchey, K., Walsh, C. A., Tong, J. K., Boxall, A. B., and Chittenden, J. P. *Phys. Plasmas* **25**(12), 122703 (2018). 173

[196] Li, C. K., Séguin, F. H., Rygg, J. R., Frenje, J. A., Manuel, M., Petrasso, R. D., Betti, R., Delettrez, J., Knauer, J. P., Marshall, F., Meyerhofer, D. D., Shvarts, D., Smalyuk, V. A., Stoeckl, C., Landen, O. L., Town, R. P., Back, C. A., and Kilkenny, J. D. *Phys. Rev. Lett.* **100**, 225001 (2008). 173

[197] Séguin, F. H., Li, C. K., Frenje, J. A., Kurebayashi, S., Petrasso, R. D., Marshall, F. J., Meyerhofer, D. D., Soures, J. M., Sangster, T. C., Stoeckl, C., Delettrez, J. A., Radha, P. B., Smalyuk, V. A., and Roberts, S. *Phys. Plasmas* **9**(8), 3558–3566 (2002). 173

[198] Frenje, J. A., Li, C. K., Séguin, F. H., Casey, D. T., Petrasso, R. D., Sangster, T. C., Betti, R., Glebov, V. Y., and Meyerhofer, D. D. *Phys. Plasmas* **16**, 042704 (2009). 173

[199] Appelbe, B., Pecover, J., and Chittenden, J. *High Energy Density Phys.* **22**, 27–36 (2017). 174

[200] Cerjan, C., Sayre, D. B., and Sepke, S. M. *Phys. Plasmas* **25**, 022705 (2018). 174