Simulating radiatively cooled pulsed-power driven systems using static mesh refinement

Nikita Chaturvedi

December 2023

Submitted in partial fulfilment of the requirements for the degree of
Doctor of Philosophy of Imperial College London

Department of Physics
Imperial College London
Prince Consort Road
London SW7 2AZ

Declaration of Originality

I hereby certify that the material in this thesis, which I submit for the degree of Doctor of Philosophy, is entirely original unless otherwise cited or acknowledged in the text.

Nikita Chaturvedi
30th December, 2023

Copyright Notice

The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a Creative Commons Attribution-Non Commercial 4.0 International Licence (CC BY-NC). Under this licence, you may copy and redistribute the material in any medium or format. You may also create and distribute modified versions of the work. This is on the condition that: you credit the author and do not use it, or any derivative works, for a commercial purpose. When reusing or sharing this work, ensure you make the licence terms clear to others by naming the licence and linking to the licence text. Where a work has been adapted, you should indicate that the work has been changed and describe those changes. Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.

Abstract

Radiative cooling is a ubiquitous process in several high energy density physics (HEDP) experiments, however numerical modelling of the final compressed state is often resolution-limited. This work presents the development of a static mesh refinement capability for use in the radiation-magnetohydrodynamics code, Chimera, in order to perform high resolution simulations of pulsed-power driven experiments. A block-based refinement strategy is used, where the grid is recursively split into blocks of equal number of cells, but at different resolutions. Each processor is allocated a single block, therefore all processors perform an equal number of computations, and load balancing is inherently achieved. The exchange of grid variables across different resolutions is performed via MPI communications. A novel algorithm is described for exchanging MHD variables in a constrained-transport solver, to maintain the divergence-free properties of the solution.

Use of this capability is demonstrated on simulations of three systems: single wire Z-pinches, X-pinches (in wire and solid configurations), and magnetic reconnection using inverse wire arrays. A systematic study of single wire experiments on a 150 kA generator is performed in 1D, 2D, and 3D at equivalent resolutions. Simulations in 1D demonstrate the need to use EoS and transport models with condensed matter corrections to perform simulations with 'cold-start' initialisation. Simulations in 2D and 3D demonstrate the effect of MHD instabilities on hotspot formation. High resolution wire X-pinch simulations in 3D are presented that, for the first time, demonstrate the 3D asymmetries that arise at the crossing point due to wire offsets. 2D axisymmetric simulations of solid X pinch experiments are presented demonstrating extreme conditions reached at the core of the crossing point. Finally, 2D and 3D simulations of radiatively-cooled magnetic reconnection experiments performed on the Z facility are presented, that demonstrate the sensitivity of the reconnection dynamics to resolution and dimensionality.

Role of the author

The development of the radiation-magnetohydrodynamics code is a result of the efforts from several existing and prior members of the Centre for Inertial Fusion Studies (CIFS) group at Imperial College London, led by Prof. J. Chittenden.

The core hydrodynamic, magnetohydrodynamic, and radiation modules, which are the focus of this work, were developed several years before the work in this thesis began. The vector potential MHD solver for use in plasma and vacuum regions was developed by Dr. C. Jennings, the reformulation into a constrained-transport solver was performed by the author. The group-resolved radiation loss, a frequently used model in the simulations presented, was developed by Dr. A. Crilly. The atomic and microphysics code SpK, and corresponding post-processing tools, were originally developed by Dr. N. Niasse and have since been further developed by Dr. A. Crilly.

The original conception and implementation of a block-structured mesh refinement capability was by Drs. C. Jennings and J. Pecover. The extended capability with arbitrary refinement levels and flexible splitting was developed by the author. The conception of the data structures used in this capability to create a tree-like hierarchy of grids was a result of discussion with Dr. N. Niasse.

Simulations presented of the pulsed-power driven magnetic reconnection experiments that compare radiation models were performed by Prof. J. Chittenden, the analysis was performed by the author, and further relevant discussion is a result of the discussion between members of the MARZ collaboration.

The work describing simulation results and analysis of solid X-pinch experiments performed on the SATURN generator has been submitted for publication by the author. It is under review at the time of writing, and is expected to be published in May 2024.

Acknowledgements

The start of the PhD seems now like many more than four years ago. A lot has happened since, and there are many who have helped me along the way.

Firstly, deep thanks go to my supervisor Prof. Jeremy Chittenden, for sharing his many MHD insights, stress-testing every single version of my several code re-writes, and his guidance all the way through.

The Plasma group at Imperial, particularly the CIFS team, has been an incredible support system. For keeping me caffeinated with a continuously supply of tea and coffee, for drawing endless cubes and arrows all over the whiteboards, for last minute thesis corrections, for belting out Troy at every possible opportunity, and for making the office an enjoyable place despite the frustrations of mesh refinement and MPI, I am immensely grateful.

Sixty miles along the M40 – I would like to sincerely thank First Light Fusion for the opportunity to pursue this PhD. The enthusiasm for physics of the people working there served as the inspiration to undertake this research. I look forward to seeing First Light's successes in the coming years and wish them the very best in their pursuits. A special mention goes to Nic, my supervisor there – for the physics, and friendship, my many thanks. And finally, to the Oxford crew, I look forward to all the times we'll be assembled in the Anchor again.

A thousand thank you's will not be enough here. To my friends and family in London – thank you for being a home away from home. To my friends and family all over the world – you are terribly missed. Dad - thank you for making me laugh when it was the last thing I wanted to do. Mum - thank you for being my rock.

Contents

1. Introduction **19**

1.1. Pulsed power and Z-pinches 19

1.2. Radiative collapse 24

1.3. The need for mesh refinement 28

1.4. Thesis overview 29

2. Theory **31**

2.1. MHD Equations 31

2.2. EOS 40

2.3. Transport 45

2.4. Radiation transport 48

3. Mesh refinement **56**

3.1. Overview 56

3.2. Review of parallelisation concepts 58

3.3. Review of refinement strategies 60

3.4. Previous implementations 64

3.5. Current implementation 71

3.6. MHD variable transfer 80

4. Single wire Z-pinches **90**

4.1. Review of dynamics 91

4.2. Experimental overview 96

4.3. 1D simulations 98

4.4. 2D simulations 105

4.5. 3D simulations 114

4.6. Concluding remarks 118

5. X-pinches **119**

5.1. Overview 119

5.2. 200kA fibre X-pinches 123

5.3. 5 MA solid X-pinches 139

5.4. Concluding remarks 157

6. Magnetic reconnection **158**

6.1. Overview 158

6.2. Exploding wire arrays for magnetic reconnection 161

6.3. 2D Simulations 165

6.4. 3D Simulations 175

6.5. Concluding remarks 177

7. Conclusions **179**

7.1. Mesh refinement capability 179

7.2. Radiatively cooled systems 180

A. Algorithms for MHD solvers **183**

List of Figures

1.1. Cross section of one of thirty six modules of the Z-generator, showing electrical components that transmit an electromagnetic pulse from outer (left) to inner (right) radius. Reproduced from Ref [1]. 20
1.2. Demonstration of Z-pinch based fusion schemes. a) Double-ended hohlraum concept, showing two opposing wire arrays encased in a gold hohlraum. Labels are (A/B): top/bottom primary hohlraums, C: capsule, D: secondary hohlraum, E: electrode structures. b) Three components of MagLIF, from left to right: axial seed magnetic field, laser preheat, compression of liner. Reproduced from Refs [2] and [3] respectively. 22
1.3. Illustration of MHD instabilities. *Left:* $m = 0$ 'sausage' instability which forms a series of necks and bulges. *Right:* $m = 1$ 'kink' instability, which forms a helical perturbation about the axis. Adapted from Ref [4]. . . . 28
1.4. Simulation of supernova eruption using Castro, an astrophysical code using the 'AMReX' library for mesh refinement. High resolution is used near the simulation origin, and along the interface of two materials that undergo mixing due to the Rayleigh-Taylor instability. Reproduced from Ref [5]. 29

2.1. Discretisation of hydrodynamic variables on a grid cell, with cell-centred primary variables ($U_{i,j,k}$) and face-centred fluxes (F, G, H). Cell faces are referred to as left/right ($i \pm 1/2$), down/up ($j \pm 1/2$), back/front ($k \pm 1/2$). 35
2.2. MHD variables stored on a staggered Yee grid, with face-centred magnetic field components and edge-centred electric field components. 36
2.3. Global magnetic field boundary conditions, used to imitate a forced current profile from a generator. Grey shaded region indicates where the magnetic field is forced, with an inverse radius dependence. The field propagates through the vacuum (white region) into the plasma (blue region), which induces a current in the plasma. 39
2.4. Phase diagram of showing classifications of states of matter as a function of density and temperature. Notable features are the warm dense matter regime where HEDP systems of interest often occur. Reproduced from [6]. 40
2.5. Demonstration of Maxwell construction procedure, where the equilibrium pressure P_{EQ} for the isotherm (black line) is obtained by matching the areas $F_A = F_B$. Modified from Ref [7]. 44

2.6. (a) Subset of isotherms (grey lines) from an aluminium FEOS table after Maxwell construction procedure shown in Fig 2.5 has been applied, dashed red line indicates solid density. (b) Same set of isotherms as in (a), but on a linear density scale and magnified close to solid density, showing a sharp drop in pressure for states slightly expanded from solid density conditions. Note, these plots are shown the inverse $\hat{x}$ axes to Fig 2.5 ($\rho \sim 1/V$). 44

2.7. (a) Comparison of Spitzer and Lee-More conductivity models against partial wave calculations. Reproduced from Ref [8]. (b) Scaling of aluminium conductivities as a function of electron temperature using LMD model, for ρ_S (red line), $\rho_S/10$ (pink line), and $\rho_S/100$ (purple line). . . 47

2.8. (a) Diagram of radiative processes that induce electronic transitions between energy levels. (b) Example spectrum of absorption coefficient of hydrogen at $T_e = 2eV, n_e = 10^{14}\text{cm}^{-3}$. Free-free contribution is continuous, shown in green. Free-bound contribution has a characteristic "saw-tooth" profile due to edges in continuum, shown in blue. Bound-bound contribution is line features (not explicitly shown). Total (shown in red) is sum of all three. Dashed vertical lines show examples spectral group boundaries. Both images are modified from Ref [9]. 51

2.9. Radiative power as a function of electron temperature using simple model (black lines) versus group-loss resolved model (red lines) for the optically thin (solid lines) and optically thick case with $R = 20$ μm (dashed lines). Produced by Dr. A. Crilly. 54

3.1. Demonstration of domain decomposition. Full spatial domain (black outline) is split into four sub-regions which are sent to four corresponding processors $P_0 - P_3$. Shaded blue cells indicates the cells on which the solver is applied for P_0, blue box around P_0 shows the surrounding ghost cells that are also required by the solver. These overlap with cells in P_1, and therefore require inter-processor communication. 58

3.2. Illustration of three refinement strategies with respect to a curved feature. 60

3.3. Demonstration of cell-based splitting the top-left corner of a 2×2 cell grid, to reach four times the original resolution. The corresponding tree structure is shown on the right, where leaf nodes are circled in black and form the simulation grid. 61

3.4. Demonstration of load balancing across processors for a uniform mesh (left), and an unstructured mesh mapped by a Z-order Morton SFC (right). 63

3.5. Demonstration of communications between processors in the $\hat{x}$-direction to populate ghost cells on $\hat{x}$-face. *Left:* 2D representation, where processor numbers are labelled in large font, underneath are respective *irx* and *iry* indices assuming a Cartesian topology. Arrows correspond to steps 1 and 4 in the series of point-to-point communication. *Right:* 3D representation, showing transfer of information from P_1 to P_0 using an extended region. Performing steps 1-4 described above for the x and y direction with this extended region ensures that the bottom right corner of P_0 is populated with information from P_4, without requiring an explicit transfer. 65

3.6. Demonstration of initial implementation of single-level mesh refinement in Chimera. Reproduced from [10]. 66

3.7. Extension of implementation in Ref [10] to 3D and arbitrary number of levels, demonstrated in two refinement configurations; (a) an extension of Fig 3.6 to 3D (b) example of three grid levels. 67

3.8. Example splitting of domain into three levels of refinement, using the extended implementation of mesh refinement. *Left:* processors (thicker black outline) with 2×2 cells shown with *irank* and *irx*/*iry* indices labelled. Cell shading shows grid level and therefore resolution level. The blue grid has lowest resolution (largest cells), and the yellow grid has the highest resolution (smallest cells). *Right:* Table describing properties of each grid, including number of processors per direction and entry/exit points for higher level grid. Note, underlying processors in $igrid = 1$ and $igrid = 2$ are not shown for clarity. 68

3.9. Distribution of data from a low resolution 'parent' processor on the left, to four higher resolution 'children' processors on the right. Communication performed using the collective `mpi_scatterv` routine. 69

3.10. *Left:* Slices through a 3D simulation of capsule implosion with an applied multi-mode perturbation using three levels of mesh refinement, where the domain boundary of each grid is outlined in white. *Right:* lineouts through the $\hat{x}-$axis for the three grid levels (orange, green, and black line for $igrid = 1, 2, 3$ respectively). 70

3.11. Illustration of the grid splitting procedure. Red box on the base grid indicates region of interest to be refined. Remaining tree hierarchy of blocks shown underneath, where shading of leaf blocks corresponds to resolution level. Traversal order to split the grid shown on the right. Data structures used to define refinement shape (red box) and blocks shown in insets. 73

3.12. Example traversal orders for navigating a tree structure. Both orders show the tree structure on the left, and the order of the nodes accessed using their respective logic on the right. 74

3.13. *Left:* example cylindrical shell refinement object, with possible configurations of intersecting and non-intersecting blocks. *Right:* logic used to determine whether a block and cylinder intersect, given their dimensions. 74
3.14. Demonstration of injection, prolongation, and restriction operations to populate ghost cells in the processor on the right (white cubes). Extended extents demonstrated in Fig 3.5 not shown here for clarity. 76
3.15. Seven possible splitting configurations using a flexible splitting approach 77
3.16. Interpolation indexing for different splitting configurations, shown for $\hat{x}$-direction ghost cells where x-indexing is not relevant. Leftmost cell is from a parent processor, right three cells are possible configurations of children cells. *Top row:* 3D view, *bottom row*: corresponding 2D $\hat{y} - \hat{z}$ view. 77
3.17. Example of mesh splitting for a 2D simulation of the MARZ platform. A total of 2962 processors were used in this variable resolution configuration. 79
3.18. Restriction procedure for magnetic field in Ref [11], for an interface between different resolution grids in the $\hat{x}$ direction. *Top:* Restriction of the B_y component, where $B'_y = 1/4 \sum_{i=0,3} b_i$ for equal area faces. Similarly applies for B_z component. *Bottom:* Restriction of the B_x component, where the low and high resolution grids share a face. The same area-weighting sum is applicable, however this is automatically achieved by replacing the coarse electric field values with an average of the finer electric field that share an edge. 81
3.19. Demonstration of asymmetric stencils for **B** and **E** on a Yee grid. . . . 82
3.20. Demonstration of **B**/**E** exchange routine for the updated magnetic field solver. P_1's ghost cells are shown in white arrows for magnetic field components, and white circles for electric field. Corner transfers are discussed in Fig 3.23. For clarity, only relevant domain cells that are involved in the transfer are shown. 83
3.21. Variable resolution version of Fig 3.20, showing restriction and prolongation operations for **B** field components. Zoomed in subset of cells shown below with labelled values. 85
3.22. *Left:* prolongation and *right:* restriction of **E** components for populating ghost cells (white circles) in P_0 (blue processor). 85
3.23. Demonstration of how components of the MHD variables are uniquely defined at each grid cell, shown here in terms of E_z. Colour of dots corresponds to which processor the value 'belongs'. Black arrows indicate direction of communication. 86
3.24. 3D representation of prolongation and restriction operations for B_y and E_z. Simultaneous B_z transfers for $E_y \sim \partial_x B_z$, and E_y transfers for $\partial_t B_z \sim \partial_x E_y$, are not shown for clarity. Recall, ghost cells belong to P_1 (green processor) for B_y transfer, and to P_0 (blue processor) for E_z. . . 87

3.25. Variations in spatial order of reconstruction schemes due to using a combination of high and low resolution cells in a stencil. 89

4.1. Cartoon of wire ablation dynamics for high current and low current wires. In both cases, a low density coronal plsama forms around a high density core. The low current wires undergo a uniform explosion, while high current wires undergo visible $m = 0$ instabilities. Reproduced from Ref [12]. 92

4.2. *Left:* Illustration of distinct coronal plasma regions formed around the wire core in a wire-array (not to scale); ablation region (dark grey), resistive region (light grey), and bulk coronal plasma accelerated towards array axis (white). *Right:* Magnitude of the integral of fluxes plotted as a function of radius for different terms of the energy equation; radiative (black), convective (green), kinetic energy (red), and thermal conduction (blue). Dashed line indicates wire core radius. Reproduced from Ref [13]. 93

4.3. Illustration of the $m = 0$ instability, showing central, uniform core and low density plasma forming neck and lobe regions. Preferential current path shown in dashed line. Reproduced from Ref [14]. 94

4.4. *Left:* Radiograph showing horizontal striations formed across a 100 μm wire, reproduced from Ref [15], *right:* radiograph showing single wire metastable liquid-vapour state resulting filaments and bubbles. Reproduced from Ref [16]. 95

4.5. Time series of Schlieren images of 15μm Al wire explosion in vacuum. Reproduced from Ref [17]. 96

4.6. *Left:* Optical streak imaging of 15 μm Al wire explosion. *Right:* 1D simulation of wire explosion on the left, showing radius of coronal plasma in solid line, current in the core, and fraction of the wire not in a condensed phase in dashed lines. Reproduced from Ref [17] and Ref [18]. 97

4.7. Evolution of 1D single wire behaviour, showing density (green) and electron temperature (blue) profiles at 5 different times. The core-corona divide forms by 7 ns, and the corona reaches peak expansion at 14 ns, after which it reverses back onto the core. Regions of constant ρ correspond to vacuum. 98

4.8. Streak image of the pressure profiles at 0.1 ns temporal resolution and 1 μm spatial resolution. Overlaid around streak image is the interface profile, shown in magenta to represent a specific set of simulation model configurations. Dashed black line indicates current pulse used (upper y-axis). 99

4.9. Comparison of equation of state models using interface profiles for TF-EoS (red line), FEOS (blue line), TF-EoS with condensed matter corrections, run in previous simulation study (grey line) from Ref [18]. Solid lines correspond to the coronal plasma-vacuum interface, dashed lines correspond to the core-corona interface. 100
4.10. Comparison of transport models using interface profiles for Spitzer (blue line) and LMD (magenta line). Gradient used to estimate rate of expansion of core regions for each model shown in dashed lines. Rates of core and coronal expansion labeled. 102
4.11. Interface profile of plasma-vacuum boundary, comparing four different radiation treatments:no loss (yellow), simple loss (magenta), group-resolved or 'improved' loss (green), and radiation transport (black). 102
4.12. Streak profiles showing electron temperature, zoomed in to $r \leq 100$ μm region, comparing late time (30-80 ns) dynamics of the wire using the improved radiation loss model (left) and the full transport case (right). Pseudo-colour range extends only to $T \leq$100 eV. Interface profiles are overlaid in green (left) and black (right) 104
4.13. Opacity spectrum for the core states in the radiation transport case at close to peak current. 105
4.14. Density streak plot comparison of uniform mesh (left) versus variable refinement mesh (right). Dashed vertical lines indicate resolution boundaries, i.e. between regions of $\Delta r = 1\mu$ m, 2μ m, 4μ m, 8μ m, and 16μ m. White line overlaid on right plot is the interface profile from the left plot. 106
4.15. Early time evolution of 2D simulations of 15 μm Al wire dynamics, showing close-up profiles at three different times. Each image shows the log mass density (top left), log current density (top right), electron temperature (bottom left), and electrical conductivity (bottom right). 109
4.16. Late time evolution of 2D simulations of 15 μm Al wire dynamics, showing full domain at three different times. Each image shows the log of mass density (left) and current density (right). 110
4.17. Axial variation of radial component of current density (J_r), axial velocity (v_x), axial electron temperature gradient ($\nabla_z T_e$), and log density (ρ) at 7 ns (top row) and 9 ns (bottom row). Axial extent of the magnified view is 240 μm. 111
4.18. *Top left:* spatial profiles of $\rho = 10^{-4}$ kg/m^3 contour as a function of time. *Bottom left:* standard deviation of profiles above as a function of time. *Right:* Fourier transforms of spatial profiles, showing the highest amplitude wavelength at each time in coloured dots. 112
4.19. Compressed state of a single pinch point along the fibre at $t = 41$ ns. *Top:* logarithm of current density showing the full wire and a close-up profile (550×550 μm^2 region). *Bottom:* pressure and electron temperature profiles of the magnified region. 113

4.20. Evolution of single wire simulated in 3D at three different times, shown in terms of log current density. Also overlaid is the mesh splitting, shown in grey lines. 115

4.21. Evolution of $m = 1$ instability across the full length of the Z-pinch, shown in terms of streamlines of the magnetic field. Colour scale indicates magnitude of magnetic field. Initial expansion is highly uniform, late time compression of corona into the core forces. *Rightmost image*: slice of density profile thresholded at $\rho > 10^2$ kg/m^3, showing high density core broken up into separate islands. 116

4.22. Magnified view of hotspot formation in 3D simulations of a 15 μm Al wire at two different times, in terms of pressure (left) and radiative loss (right). 117

5.1. Illustration of three different X-pinch configurations. The WXP (left-most) is the focus of subsequent sections. Reproduced from Ref [19]. . . 120

5.2. Different stages of X-pinch evolution, shown diagrammatically along the top row with corresponding experimental radiographs along the bottom row. From left to right: 1) formation of core-corona structure along wires, where coronal material is swept towards axis to form an axial jet 2) radial compression of neck from Lorentz force to form a central micro Z-pinch 3) soft x-ray burst from hot spot(s) 4) hard x-ray emission from electrons accelerating across gap. Adapted from Ref [20]. 121

5.3. Diagram of experimental set-up in Ref [21] showing two central X-pinches (labelled 1 and 2) and a third X-pinch (labelled 3). Direction of current from anode to cathode is shown in white arrows. 123

5.4. Radiography images of 25 μm X-pinches (top row), and 17 μm X-pinches (middle and bottom rows). Images show a subset of the full system, as demonstrated in the diagram on the bottom right. Time away from target pinch is shown above each image. Reproduced from Ref [21]. . . 124

5.5. (a) Diagrams showing dimensions of an X-pinch in 3D Cartesian geometry, in terms of wire radius r_w and half-opening angle θ_w. Elliptical cross-section in the $\hat{x} - \hat{y}$ plane shown in right image. (b) Profiles of inner and outer surface for 2D axisymmetric geometry when varying radius to maintain a constant linear mass. 127

5.6. Wire X-pinches showing the full wires (left) and the cross-section in the $\hat{x} - \hat{y}$ plane (right), for (a) 2-wire and (b) 5-wire X-pinches. 129

5.7. Effect of simulating an X-pinch in 2D axisymmetric geometry, with fixed radius but variable density along the wire to maintain constant mass, at three different times ($t = 0, 10$, and 40 ns). Shown from left to right are profiles of density, electrical conductivity, and current density. 130

5.8. Mesh splitting of a 3D X-pinch simulation, run using a total of 400 processors. Processors from grid splitting shown in black boxes, each of which corresponds to $N_x \times N_y \times N_z = 50 \times 50 \times 50$ cells. The equivalent standard uniform grid simulation at this highest resolution would require 3200 processors. Bottom left and right plots show the grid split in the $\hat{x} - \hat{z}$ and $\hat{x} - \hat{z}$ planes respectively. 131
5.9. (a) Initialisation of 17 μm wires on a 2 μm resolution grid. Cross-section of discretised wires uses $\sim 8 \times 10$ cells. (b) Staircasing effects due to discretising angled features on Cartesian grids, leading to irregular surfaces. 132
5.10. Matching the experimental load current from B-dot probes to sinusoidal profiles of different powers. Vertical red line indicates $t = 22$ ns. 132
5.11. Evolution of 3D simulations of a 2-wire X-pinch, shown by taking slices through the $\hat{y}-\hat{z}$ plane of density output, at six different times. Electrode rings shown in grey for clarity 135
5.12. Density iso-contour surfaces showing 3D effects in 2-wire X-pinches, shown at (a) 10 ns for 500 kg/m^3 and (b) 20 ns for 100 kg/m^3. Inset in (a) shows twisted rhombohedron surface where crossing-point material is not shielded by the wires, allowing it to undergo Ohmic heating and expansion. 136
5.13. Hotspot conditions in a 2-wire X-pinch at 22 ns, showing 2D slices in the $\hat{y} - \hat{z}$ plane of profiles of (a) pressure in Pa, (b) electron temperature in eV, (c) current density in A/m^2, and magnetic field in T. 138
5.14. CAD model of load hardware, showing electrodes in cyan and solid X-pinch load in grey etched. Dimensions of load are shown on the right. Note, the 35° diagnostic viewing angle. Reproduced from Ref [22]. . . . 140
5.15. Experimental current traces obtained using B-dot probes for four SATURN SXP experiments using different load materials. Overlaid are soft x-ray ($<$ 1keV) traces obtained using XRDs, corresponding to the right-hand axis. Reproduced from Ref [22]. 142
5.16. X-ray streak camera imaging of an SXP load. (a) Raw streak images of emission through a slit using ten filters, between $56 - 62$ ns. (b, top): Filter responses for six of the ten filters. (b, bottom:) spectra from a 1275 eV blackbody source and 3500 eV bremsstrahlung source. (c) Experimental intensity amplitudes from the six filters, and corresponding calculated amplitudes using the above spectra. Reproduced from Ref [22]. 143
5.17. Set up of 2D axisymmetric simulation, showing the SXP in dark grey and with grid splitting in black lines. The splitting shows processor boundaries, each of which has 50×50 cells, therefore smaller boxes correspond to smaller cell sizes. 144
5.18. Matching load current data to $\sin^n$ profiles, for varying n. Using $\sin^2$ will give excessively high current at early time compared to data. Other powers also are only a good fit over short intervals. Instead, interpolated fit of data points was used. 145

5.19. Time series of 2D axi-symmetric simulations of the solid machined X-pinch load fielded on SATURN, showing pseudo-colour of log-current density. The current is seen to diffuse into the centre, and compress on axis. 146

5.20. Late time dynamics showing peak compression and re-expansion. 147

5.21. Velocity components before and after peak compression (57.6 ns and 57.7 ns). White arrows indicate velocity direction, left half of each plot shows V_x component, right half shows V_z. 148

5.22. Comparison of states near stagnation for different radiation models . . . 149

5.23. Core and edge components of the hotspot, shown in terms of (left plot) current density and pressure, and (right plot) radiation/electron temperature on the left/right halves. The core has uniform temperatures but peaks in pressure on the axis, while the edge carries the highest current density and therefore is hotter from higher Ohmic heating rates. 151

5.24. *Top:* density (left) and electron temperature (right) profiles from simulation run with radiation transport at 57 ns. Synthetic self-emission images using 290 eV (left) and > 1 keV (right) filters, from a side-on LOS. Horizontal striations are due to 2D simulation output being azimuthally revolved to imitate 3D output. 153

5.25. Photon mean free paths $\lambda_{mfp} = 1/\kappa_\nu$ as a function of energy for two regions in the X-pinch. 153

5.26. Self-emission images from a 35° LOS, showing (a) experimental images and (b) synthetic images, for (left) lower and (right) high energy emission. 154

5.27. *Left:* digisted filter responses in coloured lines, spectrum used for calculation in black line. *Right:* cumulative sum as function of photon energy using Eq 5.14 for different filters. Note, all values shown here are normalised to the maximum value within each data set, as absolute measurements are not available. 156

6.1. Diagram of reconnection process, showing magnetic field lines in black and velocity arrows in blue. From left to right: field lines are advected into the central layer at the inflow velocity, V_{in}, undergo a rapid change of topology, and are advected out in the perpendicular direction at Alfvèn velocity, V_A. 158

6.2. Diagram showing plasmoid formation at the centre of a reconnection layer, and ejection with the outflows. These are characterised by 'O' point topology of the magnetic field. Lightly shaded arrows indicate direction of inflow, darker arrows indicate outflows. Adapted from Ref [23]. 160

6.3. Phase diagram of magnetic reconnection regimes, in terms of normalised scale length λ and Lundquist number S. Reproduced from Ref [24]. . . 160

6.4. Magnetic field topology for a standard imploding wire array (left) and exploding wire array (right), both of which have axially directed currents. Wire cores in each configuration have a local field which merge with the global field. Direction of resultant $\mathbf{j} \times \mathbf{B}$ force is indicated by blue arrows. Modified from [25] and [26] respectively. 162

6.5. Set up and results from experiments conducted on the MAGPIE facility. *Left:* diagram of two adjacent exploding wire arrays, showing direction of magnetic field and velocity. Illustration of reconnection processor highlighted along layer. *Right:* processed interferogram images shown electron density maps for carbon (top) and aluminium (bottom) experiments. Modified from Ref [27]. 163

6.6. Image of MARZ experiment load hardware, showing two wire arrays in exploding configuration with relevant dimensions labelled. Courtesy of Sandia National Laboratories. 165

6.7. Evolution of simulations of the MARZ experiment, shown at three different times (50, 150, and 250 ns from top to bottom) in terms of the logarithm of density in kg/m^3. Top image right array shows position of central conducting rod relative to array. 166

6.8. Evolution of the reconnection layer at 150 ns and 250 ns for (from left to right) magnetic field in T, current density in A/m^2, density in kg/m^3, and electron temperature in eV. 168

6.9. Variable mesh grids used to refine a) the reconnection layer only, and b) the arrays as well as the inflows entering the layer. In the latter, the $\mathbf{j} \times \mathbf{B}$ force vectors are shown to demonstrate the region of influence that is refined besides the array and layer. Inset shows the wire cores within the high resolution region, while ablated plasma enters low resolution areas. Red dashed boxes indicate regions of interest for further discussion. 169

6.10. Effects of resolving only the reconnection layer at 25 μm, at 100 ns. *Left:* images showing the layer-only region showing $T_i > 200$ eV, where image on the right has processor splitting overlaid in grey lines. *Right:* lineouts through the $\hat{y}$-axis for both simulations showing ion temperature at three different times: (top to bottom) 100, 150, and 200 ns. 170

6.11. Effect of ablation rate using wires initialised at 50μm resolution using 3×3 cells (left), and initialised at 25μm using 4×4 cells (right). Top images show segment of arrays at initialisation, where the wire spacing is not seen to be strongly affected. Bottom images compares density (in kg/m^3 of ablation streams at 50 ns for both resolutions. Grid splitting is shown overlaid in grey lines. 171

6.12. Effect of inflow and layer conditions comparing wires initialised at 50μm versus 25μm resolution. *Left:* Images of electron temperature in eV at 150 ns. Two sets of colour scales used for clarity; blue-red (0-40 eV) for inflow, and purple-yellow (50-370 eV) for central current sheet within layer. Lineouts spanning 16mm across high temperature plasmoids for 50μm (magenta line) and 25μm (red line) resolution shown on the right, in terms of y-component velocity in $\hat{y}$-direction, and electron temperature. 172

6.13. 2D simulations of reconnection dynamics comparing group-resolved radiation loss model and radiation transport model. *Left:* plots comparing experimental (yellow lines) and synthetic (green lines) B-dot probes from model options, placed 5mm away from array. Current traces (grey lines) shown on right axes. *Right:* profiles of magnetic field in T (left images) and electron temperature profiles in eV (right images) from each simulation, at 200 ns. Simulations run by Prof. J. Chittenden, experimental B-dot data courtesy of Sandia National Laboratories. 173

6.14. Effect of using radiation transport, shown at 240 ns. *Left:* profiles of electron and radiation temperature in eV. The radiation temperature shows a diffusive wave from the plasmoids, leading to a preheat of the inflow, shown for the right array in grey dashed box). *Right:* Radiation flux in $\hat{x}$-direction for two energy groups, 9-26 eV (left central image and bottom lineout) and 76-224 eV (right image, top lineout). 174

6.15. Mesh splitting for 3D simulation of wire array reconnection, showing density profile (kg/m^3) from wires at 50 ns, and processor splitting in grey outlines. Dimensions of full domain, array radius, and high resolution width provided. 175

6.16. *Top:* 3D view of ablation streams from wire arrays, showing wire cores in grey, and slices through the density profile (in kg/m^3) in the $\hat{x}-\hat{y}$ and $\hat{x}-\hat{z}$ directions. *Bottom:* slice through $\hat{x}-\hat{z}$ plane showing magnitude (top) and $\hat{x}$-component of current density (in A/m^2). 177

6.17. Top-down view of electron temperature profiles (in eV) of reconnection layer for a 2D (left image) and 3D (right image) simulation using equivalent variable resolution configuration. The plot on the right shows lineouts through $\hat{y}$-axis for both images. 178

1. Introduction

The research presented in this thesis investigates radiatively cooled HEDP (high energy density physics) systems, which undergo collapse to small regions and are therefore frequently under-resolved by standard MHD (magneto-hydrodynamics) simulations. The states produced are useful for various purposes, including the study of fundamental physics processes, and the longstanding goal of controlled fusion as an energy source. HEDP conditions, defined as pressures > 1 Mbar [28], are attained by rapidly depositing large amounts of energy in a material, compressing and heating it to extreme density and temperatures. Over the decades two primary technologies have been developed and used to achieve this: lasers and pulsed-power. The intense environments created in both are experimentally challenging to diagnose, and often require numerical modelling to understand the dynamics.

The focus of this thesis is on the latter, pulsed-power, which this chapter introduces. First, the technology is reviewed with reference to the largest machine of its kind. This is followed by a description of typical loads fielded on the generator, used in contexts including fusion and laboratory astrophysics. Next, the 'radiative collapse' process is summarised, which was originally thought to be a means to achieving HEDP or fusion-relevant conditions in Z-pinches. A complete collapse does not occur for reasons discussed in the following section. The challenges of modelling a highly convergent system are described next, which motivates the development of a mesh refinement capability.

1.1. Pulsed power and Z-pinches

Pulsed-power is a powerful and versatile technology useful for a range of applications. As the name implies, stored electrical energy is delivered to a target via an electromagnetic pulse, which is compressed in space and time to deliver immense electrical power to a load. The most commonly fielded loads are 'Z-pinches', which refer to any plasma carrying an axial current ($\mathbf{j}_z$) and azimuthal magnetic field ($\mathbf{B}_\theta$) that experiences a radial 'pinching' Lorentz force ($\mathbf{F}_r = \mathbf{j}_z \times \mathbf{B}_\theta$). This state can be assembled from a variety of pulsed-power loads, including wire arrays, liners, solid fibres, dense plasma foci, etc. These exhibit cylindrical symmetry, which naturally conforms to the topology of the compressional Lorentz force, although non-cylindrical geometries such as flyer plates have also been used to successfully generate HEDP states.

This section first describes the technology used to achieve these states, using the Z-machine as a reference. This is followed by a description of typical Z-pinch loads fielded on the machine that are used for fusion and alternative applications.

1.1.1. Machine

At the time of writing, the Z-machine at Sandia National Laboratories (SNL) is the world's largest pulsed-power facility, delivering up to 27 MA peak current and over 80 TW peak power in 100 ns [29]. It was originally commissioned in 1980 as PBFA-II, a fast proton-beam accelerator, and was refurbished in 2005 to its current design. Today, Z is used for four main purposes: stockpile stewardship, laboratory astrophysics, fundamental science, and fusion research.

A summary of the device's operation is provided here with reference to Fig 1.1, which shows the cross section of one of Z's 36 modules, arranged cylindrically around the central vacuum chamber [29]. Capacitors in Marx banks are charged over the course of minutes and discharged over $\sim \mu$s, requiring pulse compression to achieve the power quoted above. This is achieved first via water-insulated intermediate storage capacitors, which discharge over 250 ns. Laser triggered gas switches are used to transmit the pulse either synchronously across all modules to achieve faster discharges, or asynchronously in the 'long-pulse' mode, into pulse-forming lines (PFLs). These provide further temporal compression, delivering a 100 ns pulse to the vacuum-insulated region, separated from the water section via the vacuum-insulator diode-stack. Once under vacuum, magnetically-insulated transmission lines (MITLs) transmit the pulse towards the load. Strong electromagnetic stresses along transmission lines cause Ohmic heating of electrode material and field emission, leading to the desorption of contaminants along the surface that provide a lower inductance current path, and consequently a reduction

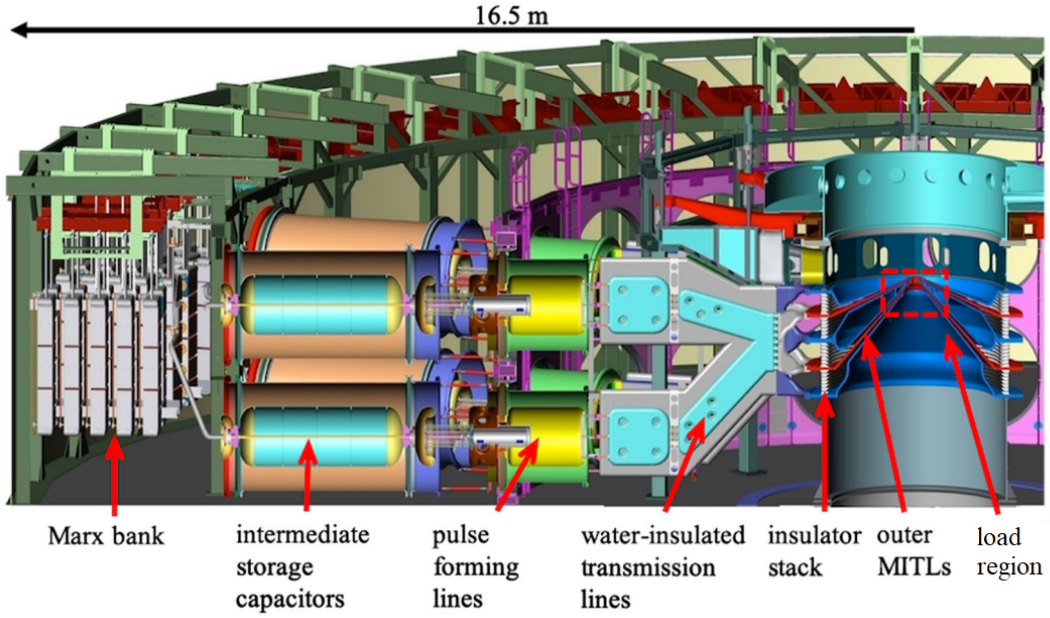


Figure 1.1.: Cross section of one of thirty six modules of the Z-generator, showing electrical components that transmit an electromagnetic pulse from outer (left) to inner (right) radius. Reproduced from Ref [1].

of power delivered to the load [1]. The MITLs act to magnetically confine electrons to the electrode via the Lorentz force, although challenges remain in achieving perfect confinement at extreme current densities. Designs for a next generation pulsed-power (NGPP) facility are currently underway, which will scale the current (and challenges) of Z to 60 MA [30].

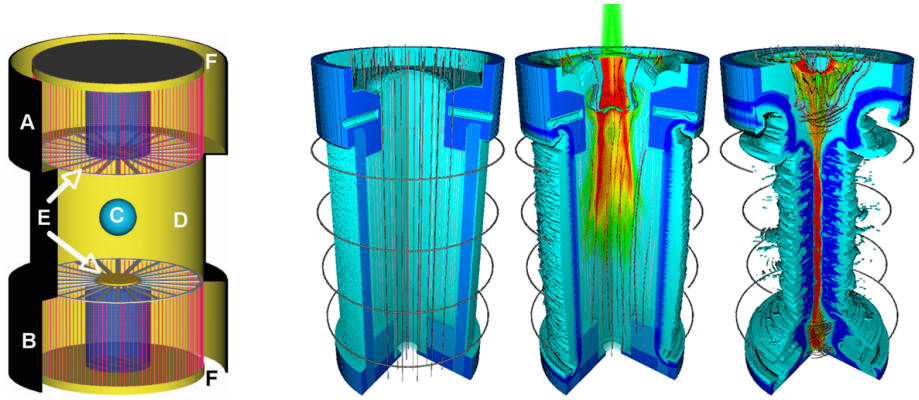
Decades of experimentation on the facility have demonstrated tungsten wire array loads to be excellent sources of soft x-rays, making them useful across all four applications listed above [31]. Here, Z's 20 MA, 100 ns current pulse is passed through a cylindrical array of wires, which are ablated and rapidly imploded towards the axis via the Lorentz force. The high velocity ablation streams collide and stagnate on axis, leading to ion viscous heating that collisionally heats the electrons enough to generate kWs of tungsten K-shell emission [32]. Overall, remarkable energy-conversion (electrical to kinetic to x-ray) efficiencies have been achieved of up to 15%, although these are limited by Rayleigh-Taylor instabilities (RTI), which are seeded from the wires and carried into the stagnated plasma column [33][31].

1.1.2. Application of Z-pinches to fusion

Thermonuclear fusion has long been considered the holy grail of clean energy, and significant effort and resource has been spent over several decades attempting to achieve it in a controlled, economical way. One approach is inertial confinement fusion (ICF), where fusion fuel is rapidly compressed and heated for a short duration, allowing sufficient fusion reactions to occur, until it is no longer confined under its own inertia and the system disassembles [34].

ICF is most notably performed in the indirect drive configuration on the National Ignition Facility (NIF) at Lawrence Livermore National Laboratories (LLNL) [35]. Here, fusion fuel, typically a deuterium-tritium mixture, is encased in a plastic capsule. The capsule is suspended in a gold casing, or 'hohlraum', which converts ∼2 MJ of incident UV laser photons from 192 NIF beams to x-rays, providing the heating 'indirectly'. The resultant radiation bath at 300 eV radiation temperature ablates the capsule, accelerating it inwards via the rocket force and generating fusion conditions in the fuel. In 2021, the facility demonstrated experiments in the burning plasma regime [36], where energy deposition of fusion products back into the fuel exceeds PdV heating, and in 2022 target energy gain was demonstrated, where more energy was produced from the reaction than deposited in, marking significant milestones in fusion research.

The conversion mechanism of electrical to laser energy, then photon energies from UV to x-ray, however, is inefficient and requires large input energies (NIF delivers ∼2 MJ to the target from 100 MJ stored in the capacitors). The conversion efficiencies associated with pulsed power is therefore an attractive alternative driver for achieving fusion-relevant conditions. This has been explored on large scale generators like the Z-facility as well as on smaller, university scale generators, using Z-pinches as loads. Three examples are described next.



(a) Double-ended hohlraum

(b) MagLIF

Figure 1.2.: Demonstration of Z-pinch based fusion schemes. a) Double-ended hohlraum concept, showing two opposing wire arrays encased in a gold hohlraum. Labels are (A/B): top/bottom primary hohlraums, C: capsule, D: secondary hohlraum, E: electrode structures. b) Three components of MagLIF, from left to right: axial seed magnetic field, laser preheat, compression of liner. Reproduced from Refs [2] and [3] respectively.

Double-ended Hohlraum

The high x-ray conversion efficiency from imploding wire arrays can be used as an alternative to lasers to perform ICF in a concept known as the double-ended hohlraum (DEH) [33]. Rather than using lasers to heat a single hohlraum, the DEH uses two tungsten wire arrays each encased in a primary hohlraum to generate an x-ray bath from either end of a secondary hohlraum that contains the fuel capsule (illustrated in Fig 1.2(a)). Experiments fielded on the Z-facility demonstrated good drive and implosion symmetry, but reached relatively low radiation temperatures for a given driver energy due to the large hohlraum volume [37]. Preliminary scaling studies using 2D radiation hydrodynamic simulations nevertheless suggest this configuration could deliver over 400 MJ yields from a 60 MA future driver such as the NGPP facility, making it a potential high-yield candidate [38].

MagLIF

Magneto-inertial fusion (MIF) schemes are magnetically-assisted fusion schemes that also rely on inertia to provide confinement. In recent years, a MIF scheme called MagLIF (Magnetised Liner Inertial Fusion) has seen considerable success on the Z facility. Here, a low-Z metallic liner containing fusion fuel is imploded, cylindrically confining and compressing the fuel to fusion conditions [39]. Liner implosion velocities are lower than that of an indirect-drive capsule, however two additional components allow the kinetic energy requirements for fusion to be relaxed. The first is an axial seed magnetic field (generated using Helmholtz coils) of up to 30 T that undergoes flux compression during the implosion to up to 100s of T, thereby suppressing thermal conduction losses from

fuel to liner. The second is laser preheating of the fuel using Z-beamlet, a TW class laser that heats the fuel to 200 eV. This serves two purposes: ionising the fuel to allow the seed field to be frozen-in, and placing it on a higher adiabat for the implosion. Neutron yields of $> 10^{13}$ (10s of kJ) using DD fuel have been demonstrated on the MagLIF platform, and similarity scaling work has demonstrated the potential to obtain yields in excess of 60 MJ on a 60 MA NGPP facility [40].

Cryogenic deuterium fibre

Both the above systems use large scale drivers and complex targets to achieve fusion conditions. However, experiments have historically been performed on smaller and simpler scales. One such example is the cryogenic deuterium fibre, for which the first experiments were performed at Los Alamos National Laboratory (LANL) [41] and then at the Naval Research Laboratory which demonstrated up to 10^{12} neutrons [42]. These high yields were initially thought to be due to a process called 'radiative collapse' (see below), where a strongly emitting Z-pinch would indefinitely collapse under the pressure of magnetic confinement, leading to fusion-relevant conditions. Several subsequent experiments were performed in the following decade on various generators in an attempt to recreate this process, including the MAGPIE facility at Imperial College London [43]. These later experiments revealed the presence of highly disruptive axial instabilities that formed in the coronal plasma, early in the current pulse. These instabilities would eventually penetrate through to the high density core, breaking it apart. Therefore instead of thermonuclear fusion occurring in uniformly compressed regions as expected, the observed yields were found to originate from energetic ion beams accelerated across the gap and undergoing beam-target processes [44].

1.1.3. Alternative applications of Z-pinches

While deuterium fibres are now considered a less viable path to fusion energy, the study of single wire Z-pinches remains useful for other purposes. Primarily, individual wires form the building blocks of more complicated, multi-wire systems such as the DEH described above. Understanding their dynamics requires an understanding of the simpler system. Wire array experiments regularly fielded on the Z-facility typically provide between 100-200 kA per wire. Therefore, generators such as IMP (150 kA, previously at Imperial College London) and XP (400 kA, at Cornell University) are ideal platforms to study wire behaviour with direct relevance to larger systems.

X-ray generation

As discussed above, imploding wire arrays are an efficient means for soft x-ray production. Another multi-wire load often used for soft x-ray production, albeit on a smaller scale, is the X-pinch. This configuration has several features of interest that are discussed later in this work. Of relevance here is the thin, axial current-carrying region, often referred to as a 'mini Z-pinch'. This region undergoes strong compression to form

hotspots of small radii from which intense x-ray bursts with short pulse durations are often observed. This behaviour makes them excellent diagnostic tools for backlighter radiography at high spatial and temporal resolution.

Laboratory astrophysics

While imploding wire arrays form Z-pinches on their axis from converging ablation streams, high density plasma flows can also be formed from the inverse configuration. Recent experiments, pioneered on the MAGPIE facility and extended to Z, use two inverse wire arrays to create diverging, magnetised outflows that interact and undergo magnetic reconnection in a dense, collisional layer. Radiatively dominated magnetic reconnection is thought to occur in extreme astrophysical environments, and lead to high energy particle outflows. At Z's high current, simulations of the reconnection layer are observed to initially go unstable to the plasmoid instability, but reach increasingly high densities and eventually undergo radiative collapse to small layer widths.

1.2. Radiative collapse

This section summarises the radiative collapse process referred to above. In 1957, Pease and Braginskii independently derived a current at which the rate of Ohmic heating in a Z-pinch is balanced by the rate of radiative cooling via bremsstrahlung emission [45]. The former (assuming Spitzer resistivities) is given by [46]:

$$P_\Omega = \eta j^2 \times \pi a^2 = \frac{I^2}{\pi a^2} \times 6.6 \times 10^{-33} Z \ln\Lambda (k_B T)^{-3/2}, \tag{1.1}$$

expressed as a power per unit length. The latter is given by:

$$P_b = 1.5 \times 10^{-29} g N_l^2 Z^3 (k_B T)^{1/2} \tag{1.2}$$

where N_l is the number density per unit length. By equating Eqs 1.1 and 1.2, and substituting the Bennett relation for a Z-pinch in pressure equilibrium:

$$(Z+1) N_l k T = \frac{\mu_0 I^2}{8\pi}, \tag{1.3}$$

one can obtain the Pease-Braginskii current:

$$I_{PB} = 2.27 \times 10^5 \sqrt{\frac{\ln\Lambda}{g}} \frac{Z+1}{Z}, \tag{1.4}$$

where $\ln\Lambda$ is the Coulomb logarithm. The well-known value of $I_{PB} = 1.4$ MA for hydrogen can be recovered by substituting $Z = 1$ and $\ln\Lambda = 10$. In order to maintain pressure balance and therefore a fixed radius for a resistive, Ohmically heated plasma, Haines [47] and Hammel [48] independently derived a current scaling profile, $I(t) \propto t^{1/3}$. However if pressure balance is not enforced and the radius is free to evolve, its rate of growth or contraction is given by [49]:

$$\gamma_r = \frac{\dot{r}}{r} = \frac{3}{4}\frac{1}{\tau_{rad}}\left[\left(\frac{I_{PB}}{I}\right)^2 - 1\right], \tag{1.5}$$

which shows that if the current exceeds I_{PB} then the pinch will contract. Based on this, the following instability mechanism is hypothesised:

1. A cylindrically symmetric Z-pinch of uniform density, initially in pressure balance, is perturbed and undergoes a small contraction.
2. Assuming purely radial motion, density is locally increased.
3. The increase in density results in an increase in radiative cooling, since $P_B \propto n^2$.
4. The decrease in temperature consequently drops the plasma thermal pressure,
5. such that the magnetic pressure exceeds the thermal pressure,
6. and the Z-pinch contracts further. The cycle is repeated indefinitely to reach immense densities and pressures.

This feedback loop is referred to as "radiative collapse".

1.2.1. Impediments

Despite numerous attempts to demonstrate this process, it has never conclusively been observed. Amongst notable experiments is the MAGPIE facility at Imperial College London, which was originally built as a high-impedance, MA-class machine to drive deuterium fibres at the Pease-Braginskii current [50]. As discussed above, this system was found to be susceptible to MHD instabilities, and evidence of collapse was not seen.

There are several simplifying assumptions underlying the above explanation that can be attributed to the absence of collapse. Ref [51] details two key mechanisms for terminating collapse: the transition to an optically thick system, and electron degeneracy effects. These are described here, along with two other mechanisms often cited in literature.

- **Optically thick regime:**
  As a plasma reaches higher densities, its emission transitions from a free-streaming to an optically thick regime. The transition is defined in terms of optical depth, defined for a given photon frequency ν over a distance s as follows [52]:

  $$\tau_\nu = \int_0^s \kappa_\nu(s) ds, \tag{1.6}$$

  where κ_ν is the extinction (absorption only in the absence of scattering) coefficient per unit frequency at given location. A description of radiation transport and atomic physics is provided in §2.4, therefore a detailed overview of extinction processes will not be repeated here. Suffice to say, electrons in partially ionised

material undergo radiative transitions, resulting in bound-free and bound-bound emission that can be more or less optically thick at different frequencies. The assumption in Eq 1.2 of pure bremsstrahlung losses is therefore not valid due to a) non-continuum processes and b) optical depth effects that act to curb the losses, making the plasma less susceptible to radiative collapse.

- **Electron degeneracy:**
  Electrons obey Fermi-Dirac statistics, with the following distribution:

$$f_{FD}(v) = \frac{1}{e^{(\epsilon-\mu)/k_B T}+1}, \tag{1.7}$$

  Under the low density/high temperature assumptions of an ideal gas, Eq 1.7 tends to a Maxwell-Boltzmann distribution. In the high density, low temperature limit, such as under condensed matter conditions, electrons are subject to the the Pauli exclusion principle, which prevents more than two fermions occupying the same quantum mechanical state. As such, the likelihood of a low-energy state being unoccupied is low, and electrons are forced into higher energy states. The pressure in these regimes as a result of quantum effects (degeneracy pressure) is then given by [53]:

$$P = \frac{\partial U}{\partial V} = \frac{2}{3}\frac{E}{V} = \frac{6}{5}\left(\frac{\pi^4}{3}\right)^{1/3}\frac{\hbar}{2m}n^{5/3}, \tag{1.8}$$

  which depends only on density and not on temperature. To give an estimate for when this is relevant, aluminium at solid density (2.7 g/cm^3) and 30 eV has a degeneracy pressure of P_{deg} =1.2 TPa. This is much higher than the equivalent ideal gas pressure $\sim$ 300 GPa, and therefore values of I_{PB} (calculated using the ideal gas expression) are an underestimate.

- **Anomalous resistivity:**
  Eq 1.1 uses a classical Spitzer version of the resistivity coefficient, η_S, which assumes simple binary collisions between electrons and ions. However, collective behaviour in certain plasma regimes can excite micro-instabilities, resulting in an anomalously high resistivity. These occur at sufficiently low line densities when the electron drift velocity is larger than the ion sound speed ($v_d^e > c_s^i$). The lower-hybrid instability (LHI) is a particular mode that, at solid fibre conditions, can grow over timescale as small as $1/\omega_{LH} \sim 10^{-12}$s, far quicker than any gross plasma motion [54]. The higher rate of Ohmic heating due to an additional anomalous component to the resistivity has been numerically observed to have adverse effects on Z-pinches in pressure balance [55]. Increased heating at the edge of the plasma causes increased expansion and difficulty in returning to equilibrium, requiring larger values of current to achieve radiative collapse, if at all.

- **Instabilities:**

Finally, one of the most egregious assumptions in the derivation is the perfectly cylindrical collapse of the pinch. Some models include the effect of axial mass losses [56], however azimuthal and axial symmetry is still assumed. Macroscopic MHD instabilities often lead to significant disruption of plasma confinement in MCF (magnetic confinement fusion) and MIF, and the two most relevant to Z-pinches are described here. Firstly, the $m = 0$ 'sausage' instability preserves azimuthal but breaks axial symmetry, and forms a series of necks and bulges across the Z-pinch length. This characteristic shape is shown in on the left of Fig 1.3. The radius at a given point varies as [57]:

$$r_{m=0}(z,t) = r_0 + \delta r \cos k_z z e^{\gamma t}, \tag{1.9}$$

where r_0 is the equilibrium radius, δr is a minor perturbation about the equilibrium, k_z is the axial wavenumber of the instability, and γ is the growth rate, which is determined by the Alvén speed, $\gamma \sim 1/\tau_A \sim r_0/v_A$. As was described above in the radiative collapse feedback loop, a small perturbation can result in rapid contraction under the magnetic field pressure. It is therefore plausible that radiative collapse could occur in the necks of a Z-pinch undergoing the $m = 0$ instability. However, the feedback loop would eventually be terminated by the bifurcation of the neck under the 'pinching' action.

The second mode is the $m = 1$ 'kink' instability, which does not assume axial or azimuthal asymmetry. This results in distortions about the axis, which compresses the field lines on the concave side and spreads out the field lines on the convex side of the instability. The characteristic helical or spiral shape is shown on the right of Fig 1.3. The radius for this mode grows as:

$$r_{m=1}(z,t) = r_0 + \delta r \cos\left(k_z z + m\theta\right) e^{\gamma t}, \tag{1.10}$$

where m is the mode number. Here, the edges of the spiral are pulled further outwards from the axis due to magnetic tension on the opposite curve, effectively breaking the uniform plasma column.

The above list of termination mechanisms is far from exhaustive, as there are many other contributions that prevent radiative collapse in a magnetically driven system. For example, the increasing inductance in a rapidly converging pinch can result in a lower current driving the system. Other macroscopic effects besides MHD instabilities such as spatial inhomogeneity in the plasma can prevent uniform compression.

Localised, brightly emitting regions are often observed from the Z-pinch plasmas described in §1.1, which compress to small radii as a the result of strong radiative cooling. However, the exact nature of termination remains unknown for several systems of interest, as the resolutions required to resolve these dynamics are often inaccessible.

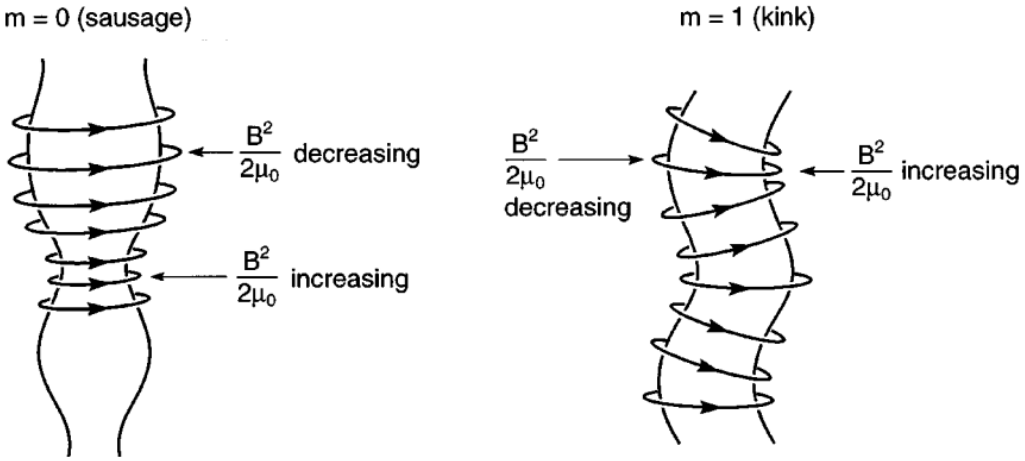


Figure 1.3.: Illustration of MHD instabilities. *Left:* $m = 0$ 'sausage' instability which forms a series of necks and bulges. *Right:* $m = 1$ 'kink' instability, which forms a helical perturbation about the axis. Adapted from Ref [4].

1.3. The need for mesh refinement

Simulations of strongly convergent systems such as the above are usually resolution-limited due to the vast range of spatial scales involved. It can be computationally intractable to resolve these fully, from initial size to compressed state. Consider a 3D simulation (using equally spaced cells) of a wire array implosion, where the initial array diameter and height are 20 mm, and the stagnated column from the snowplough implosion is compressed to 40 μm radius. To minimally resolve the column with 2 cells, the number of cells required across the full domain is at least $N = (L_x/\Delta x) \times (L_y/\Delta y) \times (L_z/\Delta z) = (20 \text{ mm} \ / \ 20 \ \mu\text{m})^3 = 10^9$ cells. Achieving even this arguably poor resolution demands significant computational resource, and simulation accuracy is often compromised instead.

One solution to the above issue is to parallelise the problem, i.e. split the domain and distribute data to separate processors, which each simultaneously apply the same solver. This approach is commonplace in scientific computing, even using smaller scale computing facilities. Consider the above example, where the grid is distributed such that processors are allocated 100^3 cells each. The simulation can then be resolved across 1000 processors, a large but not intractable number on a university-scale supercomputer. However, this still only provides 20 μm resolution. To capture the formation of small, brightly-emitting hot-spots along the stagnated column, resolutions of 5 μm or 2 μm are desirable. The former requires a 64-fold increase in number of processors, the latter requires a 1000-fold increase, only accessible at the largest supercomputing facilities.

The concept of using grids with variable resolution, referred to as 'mesh refinement' is the next solution to the 'uniform scaling' problem described above. In the late 1980s, Berger and Colella designed an algorithm to use multiple overlapping grids of varying resolution to solve the hydrodynamic equations [58]. Higher resolution grids were used

to accurately resolve discontinuous features, such as shock fronts or material interfaces, while lower resolution grids were used in smoothly varying regions containing fewer features of interest, but still necessary to capture global dynamics. An example is shown in Fig 1.4, showing the edge of a supernova explosion where materials undergo fine-scale turbulent mixing due to the Rayleigh-Taylor instability [5]. High resolution is used firstly at the core to resolve the high pressure dynamics, and secondly along the turbulent layer. The density profile is relatively smooth between these two regions, allowing it to be de-resolved and minimise the number of cells required. This technique has increasingly become a common feature in prominent hydrodynamic and MHD codes, however it requires particular attention as the grid structure underpins the entire codebase.

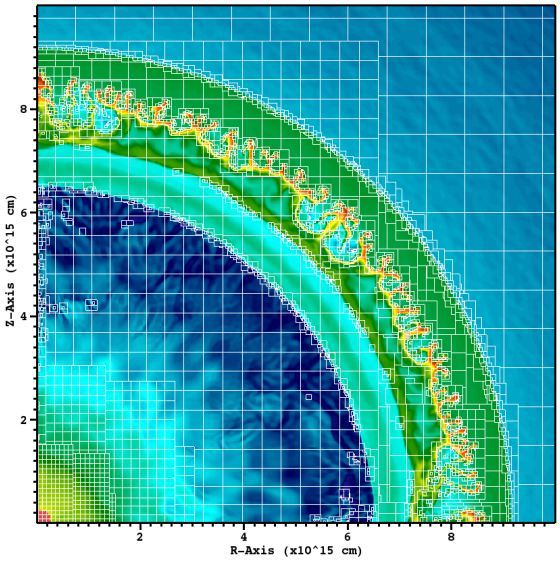


Figure 1.4.: Simulation of supernova eruption using Castro, an astrophysical code using the 'AMReX' library for mesh refinement. High resolution is used near the simulation origin, and along the interface of two materials that undergo mixing due to the Rayleigh-Taylor instability. Reproduced from Ref [5].

1.4. Thesis overview

This work describes the development of a mesh refinement capability that can be readily integrated into the radiation-MHD code Chimera, which is used to study systems undergoing strong radiative cooling, but terminating before reaching a fully collapsed state. There are several possible approaches to refinement that can be adopted to achieve this; the simplest, most robust option is chosen here. The capability is demonstrated on three pulsed-power driven systems; single wires, X-pinches, and magnetic-reconnection using exploding wire arrays. All three undergo significant cooling, and therefore undergo strong convergence to small spatial scales. Simulations of the compression of each have previously been resolution limited, this work extends the efforts to unprecedented resolutions. Studies of the single wire inform the understanding of the two other multi-wire systems. The thesis is organised as follows.

The second chapter provides an overview of the theory of relevant physical processes that occur in the three above systems, and their corresponding numerical implementation in Chimera. Specifically, these are MHD processes, radiative transfer, and equation of state closures and transport effects. The latter two are used to compare the sensitivity of systems to model options, and therefore warrant further description. Other processes such as thermal conduction and electron-ion equilibration are important but only treated using default options, therefore are not reviewed here.

The third chapter describes the development of a static mesh refinement capability, using a block-structured refinement strategy. This work extends the earlier implementation by previous developers of Chimera, which was limited to two resolution levels. An improved, more flexible implementation is described using the same underlying concept, and extended to an arbitrary number of refinement levels.

The fourth chapter focuses on the single wire system, specifically a set of aluminium wire experiments on the IMP generator. A systematic numerical study is performed of the effects of increasing dimensionality on the coronal plasma behaviour and formation of hotspots during the compression phase. The 1D simulations are used to assess the effects of different physics model options by comparing against experimental results, and ultimately justify the options used in subsequent chapters.

The fifth chapter studies the X-pinch system, where intense hotspots are produced at a central crossing point. The first ever 3D simulations are performed of a 2-wire X-pinch with an offset at the crossing point, representative of experimental conditions. Also simulated is the solid X-pinch system, which was fielded on the SATURN generator and is the highest current X-pinch system tested. This system exhibits inherent azimuthal symmetry, and therefore was modelled in 2D at higher resolution to investigate hotspot formation in detail.

The sixth chapter studies radiative cooling in a system particularly well suited to a static mesh refinement capability, pulsed-power driven magnetic reconnection. Simulations are performed of the MARZ experiments on the Z facility, which require a large domain size to drive a thin, radiatively cooled reconnection layer. The 2D simulations are used to identify regions that require high resolution. Additionally, early time results from a 3D simulation are presented that resolve the layer at 25μm, an intractable resolution without mesh refinement.

Finally, a summary of the above work is provided, as well as a review of the challenges of simulating radiatively cooling in extreme HEDP conditions at high resolution, and potential future applications of the mesh refinement capability are discussed.

2. Theory

This chapter outlines the background theory required for two purposes. Firstly, the mesh refinement capability requires knowledge of the solver to correctly transfer variables across processors, therefore discussion of the numerical techniques used in Chimera are provided. Secondly, simulations of the pulsed-power experiments described in subsequent chapter are often sensitive to the physics model used, therefore a discussion of the relevant models and their appropriate regimes is provided. The remainder of this chapter is organised as follows. The first section describes the magnetohydrodynamics equations, the second section describes equation of state models, the third section describes transport phenomena, and the final section describes radiation transport. Each provides an overview of the theoretical background, followed by a summary of its implementation in Chimera.

2.1. MHD Equations

Chimera [1] is a multi-physics code that solves the set of 3D magneto-hydrodynamics (MHD) equations in Cartesian, cylindrical, or spherical geometry. These are the set of conservation laws for mass, momentum, and energy, as well as the Maxwell equations that couple the electromagnetic field to fluid flow. The conservation equations can be derived either heuristically or using a kinetic approach, the latter is briefly demonstrated here.

The evolution of a distribution function of a species α follows the Boltzmann equation:

$$\frac{\partial f_\alpha}{\partial t} + \nabla_x \cdot (f_\alpha \mathbf{v}) + \nabla_v (f_\alpha \mathbf{a}) = \left.\frac{\partial f}{\partial t}\right|_{coll}, \tag{2.1}$$

where $f_\alpha = f_\alpha(t, \mathbf{x}, \mathbf{v})$ is the distribution function that describes the distribution of the particles across a 6-dimensional (three spatial and three velocity) phase space. The term on the right-hand side (RHS) represents collisions, and is neglected here but revisited in §2.3. Rather than solve Eq 2.1 across all dimensions, velocity moments are taken to reduce the computational expense. The zeroth, first, and second order velocity moments of f_α yield the fluid quantities: number density, velocity, and internal energy, given respectively by:

[1] Previous works refer to the code as Gorgon, however at the time of writing Gorgon is used to describe the core hydrodynamic and MHD solver, while Chimera refers to the extended capabilities.

$$n_\alpha(\mathbf{x},t) = \int f_\alpha d^3v \tag{2.2}$$

$$n_\alpha \mathbf{v}_\alpha(\mathbf{x},t) = \int f_\alpha \mathbf{v} d^3v \tag{2.3}$$

$$U_\alpha(\mathbf{x},t) = \int f_\alpha \left(\frac{1}{2} m_\alpha v^2\right) d^3v. \tag{2.4}$$

Consider a plasma comprised of electrons (subscript 'e') and a single ion species (subscript 'i'). The number conservation, or 'continuity', equations for the two particle species are obtained by taking the zeroth moment of Eq 2.1 and substituting Eq 2.2:

$$\partial_t n_e + \nabla \cdot (n_e \mathbf{v}) = 0 \tag{2.5}$$

$$\partial_t n_i + \nabla \cdot (n_i \mathbf{v}) = 0. \tag{2.6}$$

Simplifying assumptions of a single-fluid plasma can be made, where the ions are assumed to carry the bulk momentum and relative electron motion can be neglected such that:

$$\rho = \rho_e + \rho_i \approx \rho_i. \tag{2.7}$$

Eqs. 2.5 and 2.6 can therefore be multiplied by their respective specie masses, and added together to give the total plasma continuity equation:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{v}) = 0. \tag{2.8}$$

A similar procedure can be followed to add the two momenta equations to obtain a total conservation equation expressed in terms of bulk fluid quantities:

$$\frac{\partial \rho \mathbf{v}}{\partial t} + \nabla \cdot (\rho \mathbf{v}\mathbf{u}) = -\nabla(P_e + P_i) + \mathbf{j} \times \mathbf{B}, \tag{2.9}$$

where the total pressure is given by a sum of electronic and ionic components $P = P_e + P_i$, and the current density $\mathbf{j}$ is described in §2.3. The terms on the RHS represent the pressure gradient and Lorentz forces repsectively.

The same process is followed again to obtain conservation equations for the internal energy of ion and electrons. These are not added here, as Chimera evolves the temperature of each species (T_i and T_e respectively) separately.

$$\frac{\partial U_e}{\partial t} = -\nabla \cdot (U_e \mathbf{v}) - P_e \nabla \cdot \mathbf{v} + \nabla \cdot (\kappa_e \nabla T_e) + \Delta_{ei} + \eta |\mathbf{j}|^2 - \Lambda, \tag{2.10}$$

$$\frac{\partial U_i}{\partial t} = -\nabla \cdot (U_i \mathbf{v}) - P_i \nabla \cdot \mathbf{v} + \nabla \cdot (\kappa_i \nabla T_i) + \Delta_{ie}, \tag{2.11}$$

where U_α, P_α, and κ_α are the internal energy, pressure, and thermal conductivity

respectively for $\alpha = e, i$. In order, the RHS terms shared in both equations are the advection of energy, PdV compressional heating, thermal conduction, and electron-ion exchange term, where $\Delta_{ei} = -\Delta_{ie}$. The two additional terms in the electron energy equation are the Ohmic or resistive heating term due to collisions between electrons and ions, proportional to the electrical resistivity η, and a term representing the energy exchange between electrons and the radiation field. These are discussed further in §2.3 and §2.4 respectively.

The electromagnetic terms **E**, **j** and **B** are introduced via Maxwell's equations:

$$\begin{aligned} \nabla \cdot \mathbf{E} &= \frac{\sigma}{\epsilon_0} && \text{Gauss's law} \\ \nabla \cdot \mathbf{B} &= 0 && \text{No-monopoles} \\ \nabla \times \mathbf{E} &= -\frac{\partial \mathbf{B}}{\partial t} && \text{Faraday's law} \\ \nabla \times \mathbf{B} &= \mu_0 \mathbf{j} + \frac{1}{c^2}\frac{\partial \mathbf{E}}{\partial t} && \text{Ampere's law} \end{aligned} \tag{2.12}$$

where σ is the charge density, and the displacement current term has been retained in Ampere's law. These are supplemented by the generalised Ohm's law, which includes several extended-MHD terms. In this work, only the resistive version is considered:

$$\mathbf{E} = \eta \mathbf{j} - \mathbf{v} \times \mathbf{B}. \tag{2.13}$$

The induction equation can be obtained by substituting Eq 2.13 into Faraday's law, which describes the evolution of the magnetic field in a plasma:

$$\frac{\partial \mathbf{B}}{\partial t} = \nabla \times \mathbf{v} \times \mathbf{B} - \eta \nabla \times \mathbf{j} = \nabla \times \mathbf{v} \times \mathbf{B} - \frac{\eta}{\mu_0}\nabla^2 \mathbf{B}, \tag{2.14}$$

where the final substitution uses Ampere's law without the displacement current term and standard vector identities. The terms on the RHS are the advection of magnetic field with the fluid flow (frozen-in term, when $\eta \to 0$), and resistively-driven diffusion.

The techniques for numerically solving the above equations are discussed in the following sections.

2.1.1. Finite-difference methods

The set of conservation equations above are defined for continuous variables, which cannot be solved analytically and are therefore approximated to their discrete form and solved numerically. The approach to discretisation used here is referred to as finite-differencing. This is first reviewed as context for more detailed descriptions of numerical schemes.

Chimera solves the MHD and radiation transport equations on a discretised grid in one of three geometries: Cartesian, cylindrical, or spherical. The cardinal directions of each are represented as:

$$\hat{e}_1, \hat{e}_2, \hat{e}_3 = \begin{cases} \hat{x}, \hat{y}, \hat{z}, & \text{Cartesian} \\ \hat{r}, \hat{\phi}, \hat{z}, & \text{cylindrical} \\ \hat{r}, \hat{\phi}, \hat{\theta}, & \text{spherical.} \end{cases}$$

In a uniformly discretised mesh, split into $N_{e_1} \times N_{e_2} \times N_{e_3}$ cells, the spatial resolution in all three dimensions (Δe_1, Δe_2, Δe_3) is the same across each cell. The position of each cell within the grid is then indexed using i, j, k for the $\hat{e}_1$, $\hat{e}_2$, and $\hat{e}_3$-th dimension respectively. This is referred to as a logically rectilinear mesh, as the cell's spatial location within the grid can be inferred using only these three indices. Unstructured mesh geometries that do not obey this indexing are discussed in §3.1. Without the loss of generality, the following discussion assumes a Cartesian geometry.

The set of hydrodynamic equations (Eqs 2.8-2.11) are conservation laws, and can therefore be expressed in a conservative form [59]. For an arbitrary conserved variable U with spatial variation in all three dimension, this is given by:

$$\frac{\partial U(x,y,z)}{\partial t} = \frac{\partial F(U)}{\partial x} + \frac{\partial G(U)}{\partial y} + \frac{\partial H(U)}{\partial z}, \tag{2.15}$$

where F, G, and H are the fluxes in the $\hat{x}$, $\hat{y}$, and $\hat{z}$ directions respectively.

Eq 2.15 states that the change of a variable at a given spatial location is due to the contribution of fluxes entering and exiting a given 'control' volume, as shown in Fig 2.1. Therefore, in the finite-volume (FV) framework, the discrete quantity $\bar{U}$ is defined at a cell centre by taking an integral average of the continuous variable $U(x, y, z)$ over the cell volume [59]:

$$\bar{U} = \frac{1}{\Delta V} \int_z \int_y \int_x U(x,y,z) \Delta x \Delta y \Delta z. \tag{2.16}$$

The discrete value at a cell with index i, j, k is then defined as $\bar{U}^n_{i,j,k}$ at time $t = t_n$, where for simplicity the bar is dropped from further notation.

Eq 2.15 can therefore be expressed in its discrete form:

$$\frac{U^{n+1}_{i,j,k} - U^n_{i,j,k}}{\Delta t} = CFL \left[\frac{F^n_{i+1/2,j,k} - F^n_{i-1/2,j,k}}{\Delta x} + \frac{G^n_{i,j+1/2,k} - G^n_{i,j-1/2,k}}{\Delta y} + \frac{H^n_{i,j,k+1/2} - H^n_{i,j,k-1/2}}{\Delta z} \right], \tag{2.17}$$

where CFL is a multiplicative factor used to enforce the Courant-Friedrichs-Lewy condition, which requires the timestep Δt to be sufficiently small to capture the fastest signal speed ($\Delta t \leq CFL \Delta x / v_{max}$) and prevent numerical instabilities [59]. Note, the fluxes at indices $i, j, k \pm 1/2$ are defined at cell faces (see Fig 2.1). These are constructed using variables in neighbouring cells, where a larger stencil size used gives a higher accuracy estimate. The method for calculating these fluxes is described next.

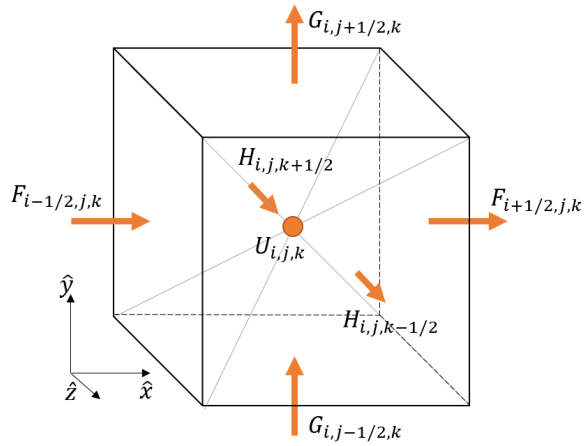


Figure 2.1.: Discretisation of hydrodynamic variables on a grid cell, with cell-centred primary variables ($U_{i,j,k}$) and face-centred fluxes (F, G, H). Cell faces are referred to as left/right ($i \pm 1/2$), down/up ($j \pm 1/2$), back/front ($k \pm 1/2$).

2.1.2. Hydrodynamic update

Chimera obtains the fluxes for each of the hydrodynamic conservation laws (Eqs 2.8-2.11) using a Lagrangian-remap scheme [60]. This splits the calculation into a Lagrangian and advection step:

1. The Lagrangian step calculates the forces on each cell face, e.g. the right face $f_{i+1/2,j,k}$, as an average of the cell-centred forces due to pressure gradients, Lorentz forces, and numerical viscosity:

$$f_{i+1/2,j,k} = (f_{i+1,j,k} + f_{i,j,k})/2, \quad \text{where} \tag{2.18}$$

$$f_{i,j,k} = \frac{P_{i+1,j,k} - P_{i-1,j,k}}{\Delta x} + \frac{j_y B_z - j_z B_y}{\Delta x} + f_{visc}, \tag{2.19}$$

   where j_y, j_z, B_y, and B_z are cell-centred quantities, and f_{visc} uses an expression for viscosity as defined in Ref [61].

2. These are applied as an upwinded (weighted in the upstream direction with respect to the flow velocity) flux in the advection step, e.g. on the right face:

$$F_{i+1/2,j,k} = \min(v_x, 0) \times U_{i+1,j,k} + \max(v_x, 0) \times U_{i,j,k}, \quad \text{where} \tag{2.20}$$

$$v_x = \frac{f_{i+1/2,j,k}}{(\rho_{i+1,j,k} - \rho_{i,j,k})/2} \times \Delta t, \tag{2.21}$$

   with an optional second-order van-Leer correction as defined by Youngs [62].

The fluxes on each cell face are then used to update $U_{i,jk}$, as per Eq 2.17.

2.1.3. MHD update

A critical aspect of any MHD solver is maintaining the divergence-free property of the magnetic field, defined in Eq 2.12 (no-monopoles). Several strategies exist to enforce this criterion [63]. For example, divergence-cleaning methods generate divergence, but apply an advection or diffusion operator to dissipate the error. Projection methods assume the magnetic field solution $\mathbf{B}$ is composed of a divergence-free solution $\mathbf{B}^*$ and an additional contribution of the form $\nabla\phi$, where ϕ is a scalar field obtained using a Poisson solver.

Chimera solves for the evolution of the magnetic field in terms of the vector potential $\mathbf{A}$, defined in term of the electric and magnetic fields as:

$$\mathbf{E} = -\frac{\partial \mathbf{A}}{\partial t} \qquad (2.22a) \qquad\qquad \nabla \times \mathbf{A} = \mathbf{B} \qquad (2.22b)$$

The vector potential is a naturally divergence-free quantity, as shown using the definition above and standard vector identities:

$$\nabla \cdot \mathbf{B} = \nabla \cdot (\nabla \times \mathbf{A}) = 0. \qquad (2.23)$$

This is enforced to numerical precision using a staggered grid, described next.

Discretisation

Unlike the discretisation of hydrodynamic equations, which stores quantities at cell centres, the MHD variables (magnetic field, electric field, vector potential, and current density) are evolved on a staggered Yee grid, which stores components at off-centre locations [64]. Specifically, magnetic field components are stored on cell faces (B_x on the $\hat{x}$-normal face, etc.), and vector potential, electric field and current density components are stored on cell edges (E_x on $\hat{x}$-edge, etc.), as illustrated in Fig 2.2.

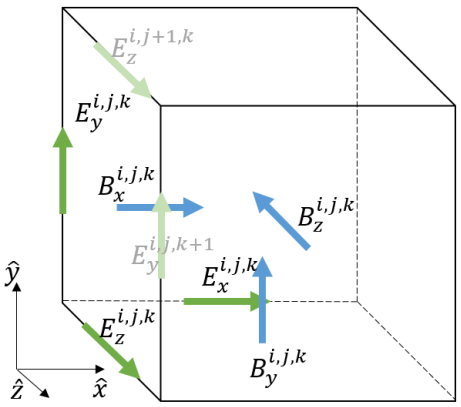


Figure 2.2.: MHD variables stored on a staggered Yee grid, with face-centred magnetic field components and edge-centred electric field components.

Using this formulation, the curl of an edge-centred variable is collocated with its

corresponding face-centred quantity. For example, the $\hat{x}$-component of the curl of $\mathbf{E}$ is collocated with B_x. This property can be shown to be divergence preserving. The finite-difference formulation for the curl of an edge-centered variable (e.g. $\mathbf{E}$) is:

$$\nabla \times \mathbf{E}|_x = \frac{\partial E_z}{\partial y} - \frac{\partial E_y}{\partial z} = \frac{E_z^{i,j+1,k} - E_z^{i,j,k}}{\Delta y} - \frac{E_y^{i,j,k+1} - E_y^{i,j,k}}{\Delta z}, \tag{2.24}$$

shown here only for the $\hat{x}$ component. Substituting this expression into the finite-difference formulation for the divergence of a face-centered variable (e.g. $\nabla \times \mathbf{E}$) gives:

$$\begin{aligned}\partial_x(\nabla \times \mathbf{E}) = &\left[\frac{(E_z^{i+1,j+1,k} - E_z^{i,j+1,k}) - (E_z^{i+1,j,k} - E_z^{i,j,k})}{\Delta x \Delta y}\right] \\ &- \left[\frac{(E_y^{i+1,j,k+1} - E_y^{i,j,k+1}) - (E_y^{i+1,j,k} - E_y^{i,j,k})}{\Delta x \Delta z}\right].\end{aligned}$$

The extension to $\hat{y}$ and $\hat{z}$ components using the appropriate Cartesian rotations is trivial, and the sum of components can be shown to be zero to numerical precision. The above expressions similarly apply to the vector potential, making this method divergence-free.

Plasma-vacuum solution

In order to smoothly transition the magnetic field between resistive plasma and vacuum regions, an expression for the vector potential is derived as follows [65]. This derivation is similar to that performed for the induction equation (Eq 2.14), however the displacement current is retained. Here, the resistive form of Ohm's law (Eq 2.13) is first substituted into Ampere's law including the displacement current (Eq 2.12) and rearranged to give:

$$\frac{\partial \mathbf{E}}{\partial t} = c^2 \nabla \times \mathbf{B} - \frac{c^2 \mu_0}{\eta} (\mathbf{E} + \mathbf{v} \times \mathbf{B}). \tag{2.25}$$

Eqs 2.22a and 2.22b are injected into Eq 2.25 to give an expression in terms of $\mathbf{A}$:

$$\frac{\partial^2 \mathbf{A}}{\partial t^2} = -c^2 \nabla \times \nabla \times \mathbf{A} - \frac{c^2 \mu_0}{\eta} \frac{\partial \mathbf{A}}{\partial t} + \frac{c^2 \mu_0}{\eta} \mathbf{v} \times \nabla \times \mathbf{A}. \tag{2.26}$$

This expression can be demonstrated to return the correct behaviour of the magnetic field in the separate vacuum and plasma limits. In the vacuum limit, $\eta \to \infty$, Eq 2.26 tends to a wave equation propagating at the speed of light, c:

$$\frac{\partial^2 \mathbf{A}}{\partial t^2} = c^2 \nabla^2 \mathbf{A}, \tag{2.27}$$

where Eq 2.23 and standard vector identities are used. In the plasma limit, $\eta \to 0$, Eq 2.26 tends to:

$$\frac{\partial \mathbf{A}}{\partial t} = \mathbf{v} \times \nabla \times \mathbf{A} = \mathbf{v} \times \mathbf{B}, \tag{2.28}$$

the curl of which can be taken to recover Eq 2.14 in the low resistivity, ideal MHD limit to show the magnetic field being frozen-in to the fluid motion.

Numerical solution

Eq 2.26 is solved in Chimera using operator splitting, i.e. the advection and resistive operators are applied separately on different timescales. The advection term (Eq 2.28) evolves the magnetic field at the fluid velocity, and therefore is solved on the same timescales as hydrodynamic variables (Δt). The diffusive terms, however, describe either resistive diffusion of the field into the plasma, or wave propagation through a vacuum, and therefore are solved on timescales (Δt_{MHD}) that capture dynamics at c.

An expression for the diffusive solution is obtained by first considering Eq 2.25 for $\mathbf{v} = 0$. This is a linear, first order ODE (ordinary differential equation), and therefore can be solved using the integration factor method to give:

$$\mathbf{E}(t + \Delta t) = \mathbf{E}(t) e^{\frac{-\mu_0 c^2}{\eta} \Delta t} + \frac{\eta}{\mu_0} \nabla \times \mathbf{B} \left(e^{\frac{-\mu_0 c^2}{\eta} \Delta t} - 1 \right). \tag{2.29}$$

The solution for $\mathbf{A}$ is obtained by injecting Eq 2.22a into Eq 2.29 to give:

$$\mathbf{A}(t + \Delta t) = \mathbf{A}(t) - \Delta t \left[\mathbf{E}(t) e^{\frac{-\mu_0 c^2}{\eta} \Delta t} + \frac{\eta}{\mu_0} \nabla \times \mathbf{B}(t) \left(e^{\frac{-\mu_0 c^2}{\eta} \Delta t} - 1 \right) \right], \tag{2.30}$$

which can be combined with Eq 2.22b to obtain $\mathbf{B}_{\text{diffusive}}(t+\Delta t)$. A leap-frog timestepping method is used to advance the solution, and is described further in Algorithm 1 in Appendix A.

Constrained-transport solver

It is highlighted here that the above method can be simplified by adopting a constrained-transport (CT) approach, which only uses $\mathbf{B}$ and $\mathbf{E}$ variables to evolve the electromagnetic field. Rather than derive an additional expression for $\mathbf{A}$, the plasma-vacuum transition is handled by directly using $\mathbf{E}$ from Eq 2.25 and Faraday's law, which is equivalent to using Eq 2.26 and Eq 2.22b. Similarly, the solution Eq 2.25 is obtained directly through Eq 2.29. The divergence-free property is similarly maintained using a staggered Yee grid, already demonstrated in terms of the required variables. As such, this method was implemented to simplify the mesh refinement procedure for exchanging MHD variable across resolution levels (described in §3.6). Some minor modifications were required to the magnetic field update algorithm, these are demonstrated in Algorithm 2 in Appendix A and compared to the above vector potential method.

MHD boundary conditions

The single and multi-wire loads modelled in §4 - 6 are experimentally driven through a series of pulsed-power components that transmit an electromagnetic (EM) pulse to

the load, as described in §1.1. The resultant EM field from the generator responds to dynamically varying loads, making this a highly coupled system. Ideally, simulations would use an integrated circuit model, which captures the current and voltage response to loads with a dynamic inductance, $\dot{L}$. However, the inductance is challenging to calculate if the load evolves to produce highly irregular topologies. Instead, the current profile from experimental B-dot (current) probes are used to approximate the total current reaching the load. In the absence of data, a sinusoidal profile approximation with assumed peak current and rise time parameters can be used instead.

In Chimera, the current is 'delivered' to the load by applying a global magnetic field boundary condition, defined using the Biot-Savart law:

$$\mathbf{B}(r) = \frac{\mu_0}{4\pi} \int \frac{I d\mathbf{l} \times \hat{\mathbf{r}}}{|\mathbf{r}|^3}, \tag{2.31}$$

which for an axial current and radial displacement vector gives the simplified form:

$$B_\theta(r > R) = \frac{\mu_0 I}{2\pi r}, \tag{2.32}$$

where R is a radial distance away from the load, such that variations of the current density profile within the central plasma can be neglected, and the approximation of an axially directed current holds. A plasma experiencing strong Ohmic heating can undergo a rapid expansion up to $r = R$. If magnetic field is advected with this expansion, the discontinuity in magnetic field between the advected and boundary value generates a large current ($\mathbf{j} \propto \nabla \times \mathbf{B}$) and can cause further Ohmic heating. This behaviour is often observed in simulations in low density regions near peak current, and can lead to premature termination of the simulation from an excessively small timestep.

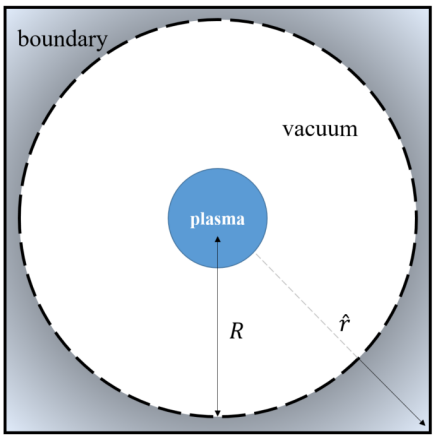


Figure 2.3.: Global magnetic field boundary conditions, used to imitate a forced current profile from a generator. Grey shaded region indicates where the magnetic field is forced, with an inverse radius dependence. The field propagates through the vacuum (white region) into the plasma (blue region), which induces a current in the plasma.

2.2. EOS

Materials in HEDP experiments can span several orders of magnitude in density and temperature. For example, a deuterium fibre fielded on the MAGPIE generator (described in §1.1.2) begins its evolution at solid density, cryogenic conditions. Within a few nanoseconds its surface is rapidly ablated to several orders of magnitude lower than solid density, and at temperatures of 100s of electronvolts Accurately modelling the material properties across this range is a considerable challenge, particularly in the warm dense matter (WDM) regime, highlighted in the phase diagram in Fig 2.4. WDM states are hotter than standard condensed matter regimes, but too dense to be described by ideal gas theory, and are challenging to diagnose therefore have little empirical data against which to validate models [66]. Inconveniently, it is through this very regime that pulsed-power systems often evolve.

Material properties across phase space are treated in simulations using an equation of state (EoS), which relates primitive state variables (density and temperature) to other thermodynamic quantities, such as pressure and sound speed. Existing models vary in complexity, from simple analytic models to comprehensive tabulated ones. Their usage depends on the region of phase space being considered, and the need for accuracy at the expense of computational resource. This section first describes underlying concepts in the construction of an EoS, then describes two models available in Chimera; Thomas-Fermi EoS (TF-EoS), and FEOS (Frankfurt Equation of State). The former is an analytic EoS, which is run inline with the solver. The latter is a tabulated EoS, which pre-calculates and stores thermodynamic variables as a function of density and temperature, and is accessed during run time.

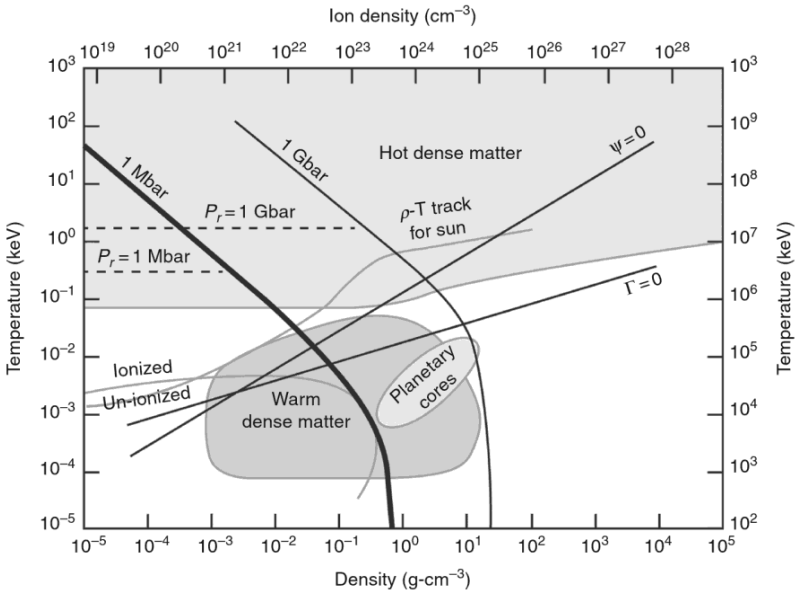


Figure 2.4.: Phase diagram of showing classifications of states of matter as a function of density and temperature. Notable features are the warm dense matter regime where HEDP systems of interest often occur. Reproduced from [6].

2.2.1. Thermodynamic relations

The hydrodynamic equations (Eqs 2.8, 2.9, and 2.10 and 2.11) were obtained by taking successively higher velocity moments of the VFP equation. The solution for each N-th moment depends on the next higher ($N+1$-th) moment, for example the solution of Eq 2.8 (conservation of mass) depends on the momentum, which itself is solved in Eq 2.9 using higher order moments. This is referred to as the closure problem.

An EoS model can be used to relate the internal energy (U, Eq 2.4) to the pressure tensor (P, considered in this work only in its scalar form), without requiring a solution to the next higher order moment of the Boltzmann equation. This is achieved via the Helmholtz free energy $F(\rho, T_e, T_i)$, which satisfies the thermodynamic relation [67]:

$$F = U - TS, \tag{2.33}$$

where S entropy and V is volume (related to density via $\rho = mN/V$) such that:

$$P = \rho^2 \left(\frac{\partial F}{\partial \rho}\right)_T \quad \text{(2.34a)} \qquad S = -\left(\frac{\partial F}{\partial T}\right)_\rho \quad \text{(2.34b)} \qquad U = F + TS, \quad \text{(2.34c)}$$

where subscripts denote the variable being held constant. The Helmholtz free energy is assumed to be additive, and can therefore be expressed in terms of its electronic and ionic contributes as [68]:

$$F(\rho, T_e, T_i) = F_i(\rho, T_i) + F_e(\rho, T_e) + F_b(\rho), \tag{2.35}$$

where F_b is a correction term used in the more advanced FEOS model to account for bonding corrections at solid conditions (see below). As a result, the derived properties (P, S, and U in Eqs 2.34(a-c)) can similarly be expressed a sum of their respective electronic and ionic components.

The electronic contribution in both models is derived from Thomas-Fermi theory, which is described in §2.2.2. However, this theory has limitations particularly in the condensed matter regime, and improvements are discussed in §2.2.3, as well as advances to the ionic contribution.

2.2.2. Thomas-Fermi

The Thomas-Fermi (TF) EoS to model electron properties was derived independently by Thomas in 1927 [69] and Fermi in 1928 [70], and is used as the foundation of several, more advanced, equations of state (including FEOS). It utilises an ion-sphere model, where electrons from multi-electron ions are treated as a degenerate fluid that sees a screened nuclear charge. As such, this model is particularly applicable to high density and/or low temperature states of matter. The description provided here follows the derivation in Ref [67].

In this treatment, an ion-sphere is defined as a sphere of electrical neutrality around an ion that contains sufficient electrons to shield the nuclear charge, with Wigner-Seitz radius:

$$r_0 = \frac{4\pi n_i}{3}^{-1/3} = \left(\frac{4\pi\rho}{3Am_p}\right)^{-1/3} \tag{2.36}$$

The Coulomb potential $V_c(r)$ that arises as a result of the central nucleus and surrounding electrons is given by:

$$-\nabla^2 V_c = 4\pi Ze\delta(r) - 4\pi e n_e(r), \tag{2.37}$$

where $n_e(r)$ is the electron number density, including bound and free electron contributions, and Z is the atomic number. The electron number density can be obtained by integrating the Fermi-Dirac distribution function (f_{FD} from Eq 1.7) over phase space to give:

$$n_e(r) = \frac{N}{V} = \frac{1}{V}\int f_{FD}(r,v)d^3rd^3v = \frac{2}{\lambda_{th}^3} I_{1/2}\left(\frac{\mu + eV_c(r)}{k_bT}\right), \tag{2.38}$$

where λ_{th} is the thermal wavelength of electrons given by $2\pi\hbar^2/m_e k_b T$, μ is the chemical potential, and $I_{1/2}$ is the Fermi integral function (I_s for $s = 1/2$).

In the self-consistent field (SCF) framework, Eqs 2.37 and 2.38 are solved simultaneously to obtain expressions for $V(r)$ and $n_e(r)$, with the boundary conditions of $V(r \to \infty) = 0$ applied to enforce charge neutrality.

The above two quantities are used to obtain ρ and T (via λ_{th}), as well as further thermodynamic quantities derived from the Helmholtz free energy, which is given by:

$$F_e = N\mu + \int \ln f_{FD}(r,v)d^3rd^3v. \tag{2.39}$$

The number of free electrons per ion using this model is given by:

$$Q = \frac{4\pi r_0^3}{3} n(r_0), \tag{2.40}$$

which can be to obtain the average ionisation, $\bar{Z}$. This model has no knowledge of shell structure, therefore $\bar{Z}$ is a continuous function of n_e. However, this simplistic approximation allows thermodynamic quantities to be straightforwardly scaled as a function of Z (atomic number), based on the relation $r_Z = r_1 Z^{1/3}$, where r_1 is the radius at $Z = 1$.

In this treatment, derived quantities tend toward the ideal gas limit for $k_BT >> \mu + eV(r)$, i.e. when interactions between neighbouring electrons are negligible relative to their thermal energy, such that:

$$P \propto \rho T. \tag{2.41}$$

In the opposite limit ($k_BT << \mu + eV(r)$) electrons behave as a degenerate gas, where

the pressure becomes independent of the temperature (Eq 1.8) such that,

$$P \propto \rho^{5/3}. \tag{2.42}$$

However, this theory has two clear deficiencies. Firstly, the continuous variation of ionisation with density does not capture effects due to discrete atomic levels (shell-structure), although these are reduced at higher density due to pressure ionisation. A larger effect is the absence of quantum correlation at solid density, which results in unphysically high pressures (GPa−MBar) at solid density. This second point is addressed in the FEOS model, described next.

2.2.3. FEOS

FEOS, or Frankurt Equation of State, is based off the QEOS (Quotidian Equation of State) model [68], and is widely used in several hydrodynamic codes for describing the properties of shocked materials. The description provided here follows Refs [7] and [68].

The model is composed of the three contributions described by Eq 2.35. The electronic contribution is obtained from the TF calculation described above. As mentioned above, one of the shortcomings of this method is unphysically large pressures of up to Mbar at solid density. in FEOS, this is accounted for via a bonding correction, of the form:

$$E_b = E_0 \left[1 - e^{b(1-(\rho_s/\rho)^{1/3}}\right], \tag{2.43}$$

where ρ_s is the solid density, and E_0 and b are values determined using the experimental bulk modulus and zero pressure condition at solid density.

The ionic contribution is a combination of several scaling laws in different regimes, as described by the Cowan model. A comprehensive description of each is provided in Ref [68], only relevant aspects are highlighted here. In the low density and high temperature limit an ideal gas scaling is applied. The Lindemann melting law is used to define the transition from solid to liquid, using:

$$T_m = \alpha \Theta_D^2 / \rho^{2/3}, \tag{2.44}$$

where Θ_D is the Debye temperature and α is a scaling constant based on the material.

The liquid-vapour transition requires a more involved treatment. Fig 2.5 shows the phase diagram of pressure versus volume at near solid densities, with two example isotherms. The first is at the critical temperature T_c, shown in grey, and the other is for $T < T_c$, shown in black. The latter shows unphysical van der Waals loops as it traverses the metastable region defined by the binodal curve (dashed blue line). FEOS eliminates these loops using a Maxwell construction procedure, which finds the equilibrium pressure (P_{EQ} in the diagram) for which the Gibbs free energy in the liquid and vapour phase are equal (equivalent to equating the areas $F_A = F_B$ in the diagram). The isotherm is then set to $\partial P/\partial \rho = 0$ at P_{EQ} within the binodal region.

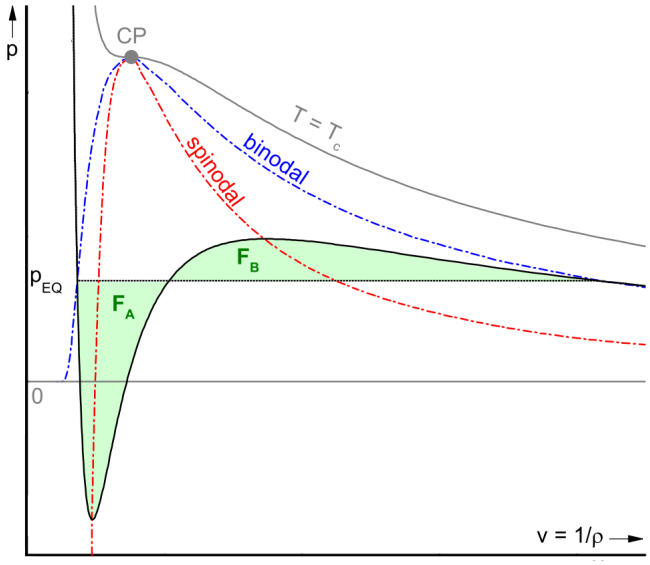
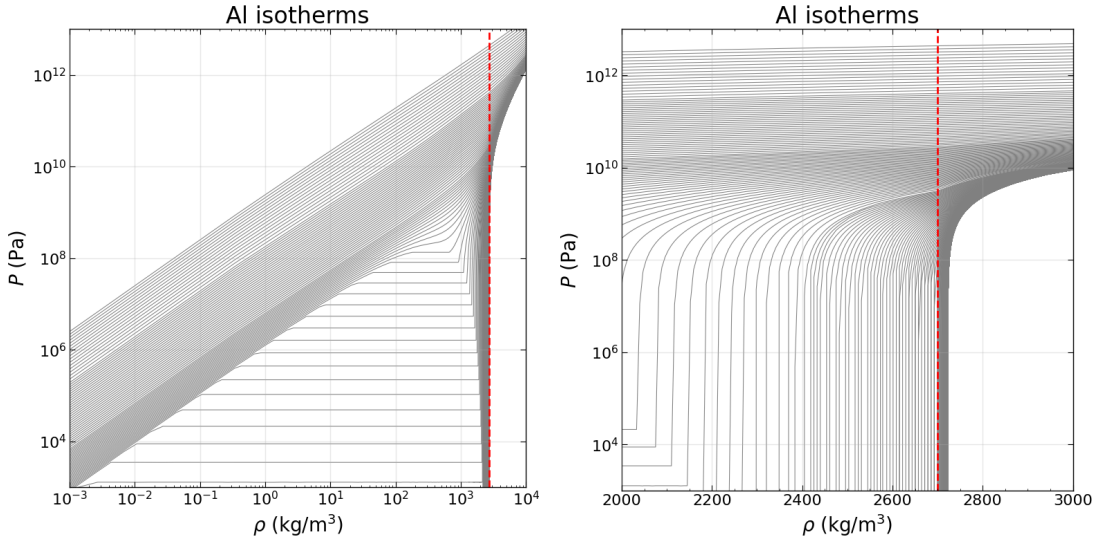


Figure 2.5.: Demonstration of Maxwell construction procedure, where the equilibrium pressure P_{EQ} for the isotherm (black line) is obtained by matching the areas $F_A = F_B$. Modified from Ref [7].



(a) Logarithmic scale (b) Linear scale

Figure 2.6.: (a) Subset of isotherms (grey lines) from an aluminium FEOS table after Maxwell construction procedure shown in Fig 2.5 has been applied, dashed red line indicates solid density. (b) Same set of isotherms as in (a), but on a linear density scale and magnified close to solid density, showing a sharp drop in pressure for states slightly expanded from solid density conditions. Note, these plots are shown the inverse $\hat{x}$ axes to Fig 2.5 ($\rho \sim 1/V$).

2.2.4. Comparison of models

The modifications in FEOS, particularly around solid density conditions, makes it a more suited model to simulations of metallic wires starting from ambient, or 'cold-start', conditions at solid density compared to TF-EoS. Using FEOS, states begin at near-zero pressures, and upon shock heating or Ohmic heating they cross isotherms to reach finite pressures. In the TF-EoS case, isotherms follow a linear scaling in $P - \rho$ space, and therefore the release behaviour using this model follows a smooth drop in temperature. In the case of FEOS, however, a small decrease in density leads to a rapid drop in temperature, giving rise oscillations in the pressure profile of materials that have been heated to near melt temperatures. Overall, however, the FEOS model uses empirically informed scaling laws in different regions of parameter space, and is therefore expected to be a higher fidelity model.

2.3. Transport

As with equations of state, accurate modelling of transport coefficients across a broad range of regimes remains a significant challenge in plasma physics. For similar reasons as above, the ideal gas regime can be easily treated, and measurements of plasma properties are straightforward to obtain at near ambient conditions. Neither is true in the intermediate, WDM regime, therefore simplifying assumptions are often made at the expense of accuracy.

Chimera provides two options for defining transport coefficients; a Spitzer-like analytical model, and the Lee-More-Desjarles (LMD) tabulated model. Magnetised transport effects are accounted for in the code using Epperlein-Haines coefficients [71], only unmagnetised coefficients are detailed here. The procedure for deriving these coefficients in each model depends on the specific regime considered. An example of one such procedure is first demonstrated in order to relate the kinetic picture described in §2.1 to the transport coefficients discussed here. This follows the description in Ref [8] for the Lee-More model, however specifics of the LMD model are described further below.

The Boltzmann equation given by Eq 2.1 can be expressed for electrons in the presence of an electric field [2] by substituting $\mathbf{a} = \mathbf{E}/m_e$ to give:

$$\frac{\partial f}{\partial t} + \mathbf{v} \cdot \nabla f - \frac{e}{m_e} \mathbf{E} \cdot \nabla_v f = \left.\frac{\partial f}{\partial t}\right|_{coll}, \tag{2.45}$$

where subscripts have been dropped for clarity.

Consider a distribution function of the form $f = f_0 + f_1$, where f_1 is the first-order perturbation about an equilibrium f_0. By substituting a simple Krook collision operator:

$$\left.\frac{\partial f}{\partial t}\right|_{coll} = -\frac{f - f_0}{\tau_c}, \tag{2.46}$$

[2] The procedure described below is performed with the inclusion of a magnetic field in Ref [68], however for the purpose of this discussion it is not necessary.

where τ_c is the electron-ion collision time, and neglecting time derivatives to only consider the steady-state, Eq 2.45 can be linearised to give

$$f = f_0 - \tau_c \left(\mathbf{v} \cdot \nabla f_0 - \frac{e}{m_e} \mathbf{E} \cdot \nabla_v f_0 \right), \tag{2.47}$$

Expressions for current density $\mathbf{j}$ and thermal flux $\mathbf{Q}$ are obtained by taking velocity moments of the perturbation. This is akin to the procedure shown in Eqs 2.2-2.4, except using the perturbation to the distribution function rather than the equilibrium value:

$$\mathbf{j} = -e \int \mathbf{v} f_1 d^3v \tag{2.48}$$

$$\mathbf{Q_e} = \int \frac{mv^2}{2} \mathbf{v} f_1 d^3v, \tag{2.49}$$

and can be related via

$$\mathbf{Q} = \mathbf{Q}_e + \frac{1}{e} \left(\frac{\mu}{T} - \frac{\partial \mu}{\partial T} \right), \tag{2.50}$$

where $\mathbf{Q}_e$ is defined in the electron centre-of-mass frame, and $\mathbf{j}$ and $\mathbf{Q}$ are defined in the ion centre-of-mass frame [71]. Substituting Eq 2.47 into Eqs 2.48-2.50 yields the following relations:

$$\mathbf{j} = \sigma(\mathbf{E} - S\nabla T) \tag{2.51}$$

$$\mathbf{Q} = TS\mathbf{j} - \kappa \nabla T, \tag{2.52}$$

where $\sigma = 1/\eta$ is the electrical conductivity, κ is the thermal conductivity, and S is a thermoelectric term. These coefficients can be expressed in terms of:

$$K_n = -\int \tau_c \frac{v^3}{3} \epsilon^n \frac{\partial f_0}{\partial \epsilon} d^3v, \qquad \text{where} \quad n = 0, 1, 2, \tag{2.53}$$

demonstrating a dependence on the electron-ion collision time τ_c. Other treatments [71] utilise a more realistic Fokker-Planck collision operator, however expressions for σ and κ can similarly be obtained in terms of τ_{ei}.

2.3.1. Spitzer

The widely used expression for electrical conductivity in a fully ionised gas as derived by Spitzer and Härm, which scales only as a function of T_e and Z, is [72]:

$$\sigma = \frac{3}{4\sqrt{2\pi}} \frac{(4\pi\epsilon_0)^2 (k_b T_e)^{3/2}}{Z e^2 m_e^{1/2} \ln \Lambda}. \tag{2.54}$$

However, the temperature scaling in the solid phase follows $d\sigma/dT > 0$, which the Spitzer model does include. In order to match experimental conductivity of metallic

wires at solid density using the Spitzer model, the initial temperature of the wire would have to be initialised hundred of eV.

2.3.2. LMD

However, the above expressions are valid for non-degenerate plasmas in the low density and high temperature limits. In an analogous manner to unphysical TF-EoS behaviour in the condensed matter regimes, use of Spitzer transport coefficients significantly underestimates electrical conductivity for solid density material. Therefore analogously to FEOS, Lee and More [8] develop a global conductivity model by combining formulas from separate models in the solid, liquid, and plasma regimes, which can be tabulated as a function of density and temperature, rather than run inline as the analytic Spitzer-like conductivity. The behaviour of this model in shown in Fig 2.7, which compares the LM scaling (excluding corrections by Desjarlais) of conductivity against electron temperature against Spitzer theory for aluminium ambient conditions. The inversion of conductivity towards the ambient conditions is visibly not captured by the simple model. Fig 2.7(b) shows the LMD behaviour at solid, one tenth, and one hundredth solid density. At lower densities (WDM regime), the material is treated as an insulator, and at high temperatures recovers a Spitzer scaling.

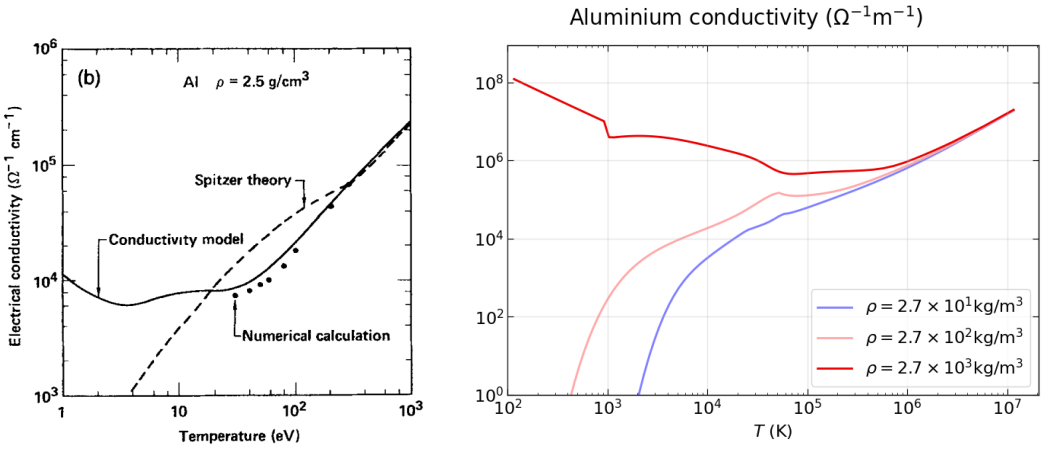


Figure 2.7.: (a) Comparison of Spitzer and Lee-More conductivity models against partial wave calculations. Reproduced from Ref [8]. (b) Scaling of aluminium conductivities as a function of electron temperature using LMD model, for ρ_S (red line), $\rho_S/10$ (pink line), and $\rho_S/100$ (purple line).

Although use of this model is a considerable improvement when simulating experiments from cold-start conditions, there are possible discrepancies between the two tabulated EoS and transport models (FEOS and LMD). Both cover similar density and temperature ranges, however transition between phases e.g. the melt line are not necessarily matched across them, which can lead to inconsistencies when they are used simultaneous. Work is currently underway to unify the scaling laws that underlie both.

2.4. Radiation transport

This section describes the framework used to model high energy radiation as it interacts with HEDP matter. The interaction of matter with longer wavelength radiation is described within the MHD framework described above in §2.1, and leads to macroscopic, collective behaviour. The interaction with smaller wavelength radiation, as described in this section, occurs on a microscopic scale with individual electrons, particularly in the x-ray range where photon energies are comparable to atomic energy transitions.

As such, radiation is typically treated in a hydrodynamic/MHD code as a collection of rays under the assumption of geometric optics:

$$dAd\Omega >> d\lambda^2,$$
$$dt >> 1/d\nu.$$

The above expressions require the radiation wavelength to be much smaller than the system's characteristic length scales, and its frequency to be much higher than the inverse of characteristic time scales. Wave-like behaviour such as diffraction and interference are therefore neglected, and rays are assumed to travel in straight lines until they are absorbed, re-emitted, or refracted (not considered here).

2.4.1. Equations of radiative transfer

The radiative energy that travels through a region of unit area dA, subtending a solid angle $d\Omega$, with unit frequency $d\nu$, and during unit time dt is defined as:

$$dE = I_\nu dAd\Omega d\nu dt, \tag{2.55}$$

where I_ν is the specific intensity with a spectral dependence. This specific intensity is related to the photon distribution function via:

$$I_\nu = \frac{2h\nu^3}{c^2} f_\nu. \tag{2.56}$$

For a photon distribution that obeys Bose-Einstein statistics,

$$f_\nu = f_{BE} \equiv \frac{1}{e^{h\nu/kT} - 1} \tag{2.57}$$

the Planck function B_ν is recovered:

$$I_\nu = B_\nu \equiv \frac{2h\nu^3}{c^2} \frac{1}{e^{h\nu/kT} - 1}. \tag{2.58}$$

The evolution of intensity is then described by the radiative transfer equation:

$$\frac{\partial I_\nu}{\partial s} = \frac{1}{c}\frac{\partial I_\nu}{\partial t} + \Omega \cdot \nabla I_\nu = -\alpha_\nu I_\nu + \eta_\nu, \tag{2.59}$$

where the LHS is the change of intensity with path length s, expanded into its temporal and spatial dependence. The RHS terms describe radiation absorption and emission; α_ν is the absorption coefficient and η_ν is the emissivity, which are related via the source function S_ν:

$$S_\nu \equiv \frac{\eta_\nu}{\alpha_\nu}. \tag{2.60}$$

To fully solve Eq 2.59 would require solving over a 7-dimensional space [3], which is computationally intractable. To reduce dimensionality, moments of Eq 2.59 can be taken with respect to solid angle. The following quantities are first defined:

$$\text{0th moment: } j_\nu = \frac{1}{4\pi}\int I_\nu d\Omega \tag{2.61}$$

$$\text{1st moment: } \mathbf{F}_\nu = \frac{1}{4\pi}\int \hat{\mathbf{n}} I_\nu d\Omega \tag{2.62}$$

$$\text{2nd moment: } \mathbf{P}_\nu = \frac{1}{4\pi}\int \hat{\mathbf{n}}\hat{\mathbf{n}} I_\nu d\Omega \tag{2.63}$$

where j_ν is the angle-averaged specific intensity, $\mathbf{F}_\nu$ is the radiation flux, and $\mathbf{P}_\nu$ is the radiation pressure tensor. Note, the energy density of the radiation field is then defined as

$$E_\nu = \frac{4\pi}{c} j_\nu = \frac{1}{c}\int I_\nu d\Omega. \tag{2.64}$$

Taken the 0th and 1st moments of Eq 2.59 and substituting Eqs 2.61-2.64 yields:

$$\frac{\partial E_\nu}{\partial t} + \nabla \cdot \mathbf{F}_\nu = -\alpha_\nu c E_\nu + 4\pi\eta_\nu \tag{2.65}$$

$$\frac{\partial \mathbf{F}_\nu}{\partial t} + \nabla \cdot \mathbf{P}_\nu = -\kappa_\nu \mathbf{F}_\nu. \tag{2.66}$$

Analogously to the conservation laws described in §2.1, Eqs 2.65 and 2.66 represent the conservation of radiation energy density, and the conservation of radiation momentum density [73]. A closure problem is evident, similar to taking velocity moments in hydrodynamics, whereby the flux of one quantity depends on its higher-order moment. A simple way to resolve this closure is to employ the Eddington approximation, which assumes that the radiation field is isotropic [73]:

$$\mathbf{P}_\nu = \frac{1}{3} E_\nu \mathbf{I}. \tag{2.67}$$

This approximation is valid in optically thick regimes, where radiation propagates diffusively. However, in the optically thin limit an incorrect speed of propagation is obtained. Substituting Eq 2.67 into Eq 2.66 and setting $\kappa_\nu = 0$ for a vacuum yields a

[3] 3 spatial, 2 angular, 1 temporal, 1 radiation frequency

propagation speed of $c/\sqrt{3}$. The $P_{1/3}$ approximation addresses this issue by introducing an additional factor of 1/3 to the time derivative term of Eq 2.66:

$$\frac{1}{3}\frac{\partial \mathbf{F}_\nu}{\partial t} + \frac{c}{3}\nabla E = -\kappa_\nu \mathbf{F}_\nu.$$

A numerical solution to this approximation is discussed in §2.4.3.

Integrating out the angular dependence as demonstrated in Eqs 2.61-2.63 reduces the phase space to 5 dimensions, one more than the hydrodynamic equations. This final variable is photon frequency (or photon energy, terms are used interchangeably in this work), on which the absorption and emission coefficients strongly depend. The next section overviews this dependence, and ways to reduce the computational expense of this extra dimension through averaging techniques.

2.4.2. Atomic physics

Atomic processes

Atomic interactions with radiation can be described in terms of three species: photons, free electrons, and bound electrons. By considering the interactions between each of them a full description of absorption and emission coefficients can be obtained. These fall into one of three categories:

1. free-free (ff),
2. free-bound (fb),
3. and bound-bound (bb),

The first deals with freely moving charges, which are accelerated or decelerated either due to other charged particles or external fields. In a plasma, this is primarily from electrons accelerating/decelerating due to interactions with ions. The deceleration results in bremsstrahlung (*'braking radiation'*) emission, where a photon is emitted with an energy corresponding to the loss of electron kinetic energy. Bremsstrahlung (and its inverse process) are continuum processes. The cross-section therefore has a simple functional form [9]:

$$\alpha_\nu^{ff} \propto \frac{1}{\nu^3}\frac{Z^2 n_e n_i}{\sqrt{T_e}} \tag{2.68}$$

The second describes radiative recombination (and its reverse process, photo-ionisation), where a free electron is captured into a bound state and releases a photon with energy $h\nu = \epsilon - E_n$, where ϵ is the electron kinetic energy and $|E_n|$ is the ionisation energy of an electron in the n-th shell. This is a continuum process above a threshold energy $h\nu_0$, corresponding to the ionisation energy of quantum levels, giving rise to a 'saw-tooth' profile. Its absorption cross-section is varies as:

$$\alpha_\nu^{fb} \propto n_e \left(\frac{\nu_0}{\nu}\right)^3 \tag{2.69}$$

Finally, bound-bound interactions consider transitions of bound electrons across energy levels, due to collisionally and radiatively driven excitation and de-excitation processes. These are resonant processes that give rise to enhanced absorption and excitation line features in the spectrum.

$$\alpha_\nu^{bb} = n_i \sigma_{i,j}(\nu) \left[1 - \frac{g_i n_j}{g_j n_i} \right], \tag{2.70}$$

where i, j are representative electron levels, $g_{i,j}$ their respective statistical weights, and $n_{i,j}$ are the electron populations in each level, and σ_{ij} is the transition cross section.

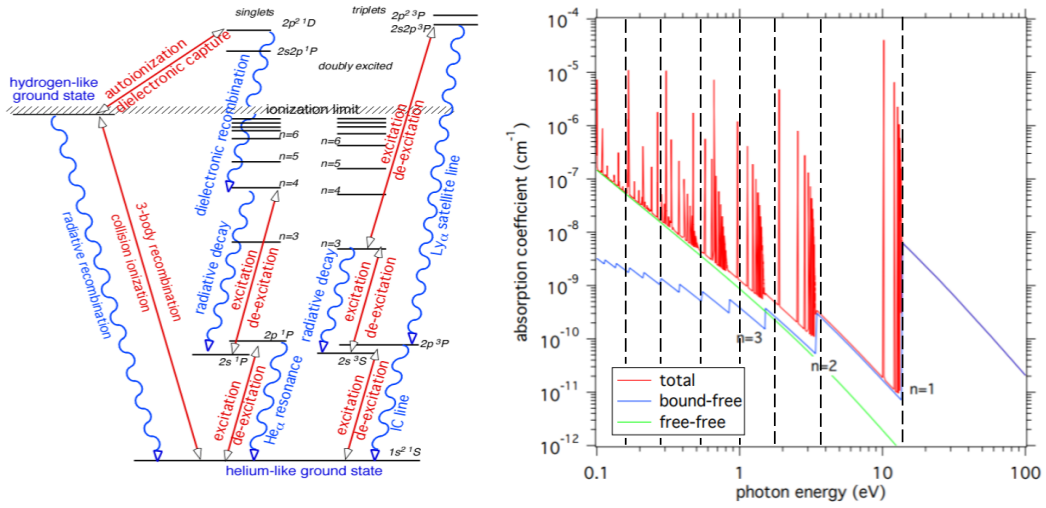
The total absorption cross section is a combination of the above, summed over quantum energy level combinations:

$$\alpha_\nu = \sum_{i,j} \left(\alpha_\nu^{fb} + \alpha_\nu^{bb} \right) + \alpha_\nu^{ff}, \tag{2.71}$$

and is often instead expressed in terms of the opacity (κ_ν),

$$\alpha_\nu = \kappa_\nu \rho. \tag{2.72}$$

An example spectrum is shown in Fig 2.8(b), that shows the total contribution (red line) as a sum of the three components. The code used to calculate spectrally-resolved opacity and emissivity data for use in Chimera is described next.



(a) (b)

Figure 2.8.: (a) Diagram of radiative processes that induce electronic transitions between energy levels. (b) Example spectrum of absorption coefficient of hydrogen at $T_e = 2eV, n_e = 10^{14}\text{cm}^{-3}$. Free-free contribution is continuous, shown in green. Free-bound contribution has a characteristic "saw-tooth" profile due to edges in continuum, shown in blue. Bound-bound contribution is line features (not explicitly shown). Total (shown in red) is sum of all three. Dashed vertical lines show examples spectral group boundaries. Both images are modified from Ref [9].

Calculation of opacities and emissivities

Chimera uses opacity and emissivity data provided by the atomic and microphysics code SpK, developed at Imperial College London [74] [75]. The expression in Eq 2.70 requires knowledge of the populations of bound/free electron species at each energy level, which can be solved using a series of coupled rate equations to account for collisional and radiative transitions, via the Saha-Boltzmann equation under the assumption of local thermodynamic equilibrium (LTE). SpK performs DCA (detailed configuration accounting) calculations to obtain population levels, using energy levels from an SHM-l (Screened Hydrogenic Model with l-splitting) and supplemented by the NIST database. A modification of the Saha model is used to perform calculations under either LTE, or CRE (collisional radiative equilibrium) conditions which are applicable to low density, mid-Z materials [75].

Averaging

The opacity and emissivity spectra produced by SpK, such as in Fig 2.8(b), are generated for a range of densities, temperatures, and average ionisations. To additionally store values at high spectral resolution is memory intensive therefore averaging techniques are required. The full spectral calculation is performed over $\sim 10^4$ points; these are then averaged and stored for a smaller number of bins or groups N_g, which is determined by the user. For tables used as input to a radiation solver, the number of groups has implications for memory and additional subcycles (run time). Typical values used are approximately $N_g = 32 - 64$, depending on simulation size versus need for accuracy. For post-processing calculations, for example to obtain 2D synthetic radiograph images, the number of groups is typically $N_g \sim \mathcal{O}(100)$.

2.4.3. Numerical implementations

This section describes the numerical implementations in Chimera to represent radiative processes, in order of increasing simplicity. The sensitivity of radiatively cooled systems to different radiation treatments are discussed in subsequent chapters. These are no loss, simple loss, group-resolved (or 'improved' loss), and finally radiation transport. The first requires no discussion, the remaining three are summarised here.

Simple loss

This treatment uses an expression for the power loss from electrons due to bremsstrahlung (ff) and radiative recombination processes (fb) [76]:

$$P_{ff+fb} = 1.69 \times 10^{-38} n_e T_e^{1/2} f_{bb} \sum Z^2 n_i(Z) \left(1 + \frac{E_\infty}{T_e}\right) \qquad [\mathrm{W/m^3}], \tag{2.73}$$

where f_{bb} is a multiplicative factor used to account for bound-bound processes and

E_∞ is the ionisation energy. This expression is substituted into the electron energy equation (Eq 2.10) as a sink term, $\Lambda = P_{ff+fb}$. This approximation is valid for highly ionised regimes, where continuum processes are the dominant contribution to radiation. However, line emission is important in partially-ionised conditions, such as in simulations described in subsequent chapters, requiring an improved model (described next). The above expression is applied as a local loss (n_e, n_i, and T_e are cell values), and therefore has no information about its neighbours. It is typically applied under an optical thin assumption, however simulations results presented in subsequent chapters use an additional attenuation factor to scale the emission. This is described in a more advanced implementation next.

Group-resolved loss

A model with intermediate complexity between a simplified loss term and full radiation transport (described next) is shown here, where the loss is calculated in a group-resolved manner to account for spectral variations using opacity (κ_ν) and emissivity (j_ν) data from SpK. Each cell is treated as a spherical emitter, with radius $R \sim \Delta x$, and the surface area (S) and volume (V) appropriately calculated. The loss is expressed as [77]:

$$P_\nu = \frac{3}{4R} \int \frac{4\pi j_\nu}{\kappa_\nu} \left[1 + \frac{2}{\tau_\nu^2}((1+\tau_\nu)e^{-\tau_\nu} - 1) \right], \tag{2.74}$$

where the term in square brackets is the probability of escape tuning factor, which depends on the optical depth $\tau_\nu = 2\kappa_\nu R$. In the limit of large optical depth, the loss tends to:

$$\lim_{\tau_\nu \to \infty} P_\nu = \frac{3}{4R} \int \frac{4\pi j_\nu}{\kappa_\nu} d\nu = \frac{S}{V} \sigma T^4, \tag{2.75}$$

where S and V are the surface area and volume of the corresponding sphere, giving the expression for a surface blackbody emitter. In the optically thin limit, the loss becomes the total (spectrally averaged) emissivity J:

$$\lim_{\tau_\nu \to \infty} P_\nu = \int 4\pi j_\nu d\nu = J. \tag{2.76}$$

A demonstration of including bound-bound effects in the power loss calculation is shown in Fig 2.9, for aluminium at one hundredth solid density. The plot shows the power calculated as a function of electron temperatures (between 1 eV and 1 keV), and integrated over the spectral range 1 eV to 1 keV. The calculation including α_ν^{bb} in Eq 2.71 are shown in red lines, and without the term in black lines. Solid lines represent the optically thin limit (Eq 2.76), and dashed lines use Eq 2.74 with $R = 20$ μm. At low temperatures bound-bound effects are negligible, however at high temperatures the calculations begin to strongly diverge. At temperatures $T_e > 100$ eV, frequently observed in the simulations presented, the aluminium K-shell edge becomes a dominant source and neglecting line emission leads to order of magnitude errors.

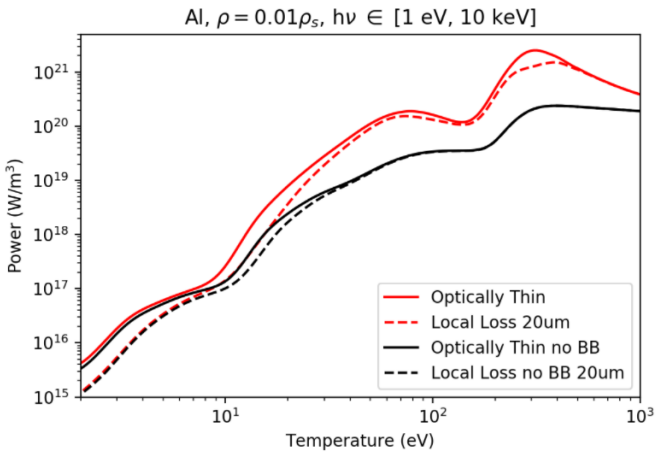


Figure 2.9.: Radiative power as a function of electron temperature using simple model (black lines) versus group-loss resolved model (red lines) for the optically thin (solid lines) and optically thick case with $R = 20$ μm (dashed lines). Produced by Dr. A. Crilly.

Despite the inclusion of line radiation, the group-resolved treatment is still a local loss model, and therefore limited in its capability. The treatment of radiative transfer is described next.

Radiation transport

Rather than using approximations as in the methods above, the zeroth and first angular moment of the equation of radiative transfer can be solved using the $P_{1/3}$ closure. The latter derivation of the latter was described in §2.4.1, and the result is repeated here for clarity.

$$\frac{1}{3}\frac{\partial \mathbf{F}_\nu}{\partial t} + \frac{c}{3}\nabla E = -\kappa_\nu \mathbf{F}_\nu. \tag{2.77}$$

Using the same integration-factor method applied to the electric field solution (obtaining Eq 2.29 from Eq 2.25), the solution to the first-order ODE in Eq 2.77 can be shown to give [65]:

$$\mathbf{F}^{n+1} = \mathbf{F}^n e^{-3c\kappa_\nu \Delta t} + \frac{c\nabla E^n}{3\kappa_\nu}\left(e^{-3c\kappa_\nu \Delta t} - 1\right). \tag{2.78}$$

In the optically thin limit ($\kappa \to 0$) this expression tends to:

$$\mathbf{F}^{n+1} = \mathbf{F}^n - \Delta t c^2 \nabla E^n, \tag{2.79}$$

which describes the free-streaming behaviour of photons, and in the optically thick, diffusive limit ($\kappa \to \infty$) it tends to:

$$\mathbf{F}^{n+1} = -\frac{c\nabla E^n}{3\kappa_\nu}, \tag{2.80}$$

which leads to a Fick's law type diffusion scheme. A leapfrog time integration scheme is performed, as with the magnetic field solver, to update the radiation energy and momentum densities.

The evolution of these quantities occurs at the speed of light c, significantly faster than any hydrodynamic motion. As a result, the radiation can be emitted or absorbed from a given cell, but also evolve 'non-locally', i.e. propagate further than hydrodynamics permits on a sub-cycled time frame. As an example, in the case of a strongly shocked systems, the photons emitted from the shock-heated material can travel faster than the shock front, heating the unshocked material ahead of it and forming an ionisation front. This preheated region is referred to as a radiative precursor, and can cause the smoothing of shock front. Loss models such as those described above are unable to capture these non-local effects, and therefore may not correctly capture dynamics.

However, the requirement of resolving dynamics at c requires sub-cycling, as is performed in the diffusive update of the MHD solver, although savings can be made by using a 'reduced' speed of light. This requirement, in conjunction with having to solve Eq 2.78 over an additional dimension for spectral resolution, makes this method of radiation treatment considerably more computationally expensive than any other physics operator used in this work.

3. Mesh refinement

3.1. Overview

The motivation for upgrading a code to include mesh refinement has been made in §1.3, which this chapter describes the implementation of. Details of the grid [1] structure underpin any numerical solver and therefore requires careful consideration. Since its conception in the 1980s by Berger and Collela [58] mesh refinement has been implemented in numerous ways, each of them bespoke to the solver and the problem being modelled. Consider two different examples, firstly simulations of the boundary layer formed along a static aerofoil immersed in fluid flow. This is a thin, turbulent layer along the aerofoil's curved boundary, which strongly affects the fluid dynamics close to the aerofoil but has little influence elsewhere. This problem has a different set of refinement requirements to another example of large scale cosmological simulation of galaxy dynamics. In this case a large domain is simulated that is primarily sparse except for two galaxies that move rapidly across the domain under the effect of their mutual gravity. For the former, a mesh will only need to be generated once as the location of the boundary layer is well known, but will require refinement around an irregular contour. For the latter, each galaxy can be modelled using regions of nested Cartesian grids in a larger base grid, however the mesh will need to be redefined dynamically with the galaxies' movement.

The above examples gives rise to two classifications of mesh refinement. The first is whether the mesh is static or dynamic. In the case of static mesh refinement (SMR) the grid is generated only once on initialisation, in adaptive mesh refinement (AMR) it is regenerated during simulation run time. A clear benefit of AMR is the ability to dynamically track features of interest, for example shock fronts or material interfaces, and to de-resolve the grid where variables are smoothly varying. However, there are several challenges associated with the flexibility of adaptive refinement. One example is the computational overhead of repeatedly generating grid topologies every time the variables have sufficiently evolved. Another is synchronising grid data across resolution levels for these continuously varying grid topologies. In these respects static refinement is less demanding and simpler to implement, although at the expense of flexibility. The second classification is the refinement strategy which determines the resultant grid topology. While naming convention varies in literature, these are identified here as one of the following; cell-based, patch-based, or block-based. The first refines the mesh by recursively splitting each cell, while the second and third form series of nested, logically

[1] The terms 'mesh' and 'grid' will be used interchangeably during this work.

rectilinear grids. The former allows better flexibility and is better suited to the curved aerofoil example above, while the latter two offer other advantages.

This chapter describes the development and implementation of an SMR library using a block-based refinement strategy in Chimera. An earlier version of Drs J. Pecover and C. Jennings was initially implemented as a proof-of-concept with the intention of integrating the MHD-transfer algorithm into an existing AMR library. Instead, the splitting strategy was used to develop a separate in-house SMR code, and the current implementation is a significant advancement of the initial principle. Further upgrades of this improved capability to an adaptive and/or cell-based approach were deliberately considered out of the scope of this work for two reasons;

- Several open-source libraries for this purpose already exist which have been developed and optimised over many years, and can be readily integrated.
- The existing implementation was shown to robustly model a large number of relevant problems. The systems considered in subsequent chapters all undergo radiative cooling at known spatial locations, obviating the need for an adaptive capability.

The remaining discussion in this chapter is organised as follows. First, §3.2 an introduction to standard concepts in parallel computing. Next, a review of the three aforementioned refinement strategies is provided in further detail in §3.3. The relevant data structures are described as these reveal the advantages and disadvantages of each strategy. With an understanding of established methods in literature, a review of refinement strategies of older implementations in Chimera is then given in §3.4. These provide the context for a more extensive description of the current mesh refinement implementation, provided in §3.5. §3.6 explicitly deals with the transfer of MHD variables, including an overview of the algorithm developed in this work.

3.2. Review of parallelisation concepts

Before describing mesh refinement strategies, an overview of the standard parallelisation techniques used in high performance (scientific) computing is provided. Concepts regarding communication across processors are pertinent to both uniform and variable mesh grids. As performance improvements of a single processing unit reaches its physical limits, the time taken to run a program can alternatively be reduced by performing the same set of instructions on several subsets of the data simultaneously. The strategy most frequently used for dividing the data is called domain decomposition (as opposed to, for example, loop-level parallelism), whereby the full spatial domain of a simulation is divided into several sub-domains and distributed to each computational processor. This is illustrated in Fig 3.1, where a domain of 16×16 cells is split into four subdomains of 4×4 cells, each corresponding to one of four processors $P_0 - P_3$.

The method for communicating data across processors depends on the computational architecture on which the solver is run. These fall under two types of classifications; shared memory and distributed memory architectures [78] [79]. As the name suggests, processors in the former have shared access to a global memory space, and therefore do not need explicit instructions for communication. However, this architecture requires an operating system (OS) to prevent conflicts when processors simultaneously contend for access to the same memory address. The complexity associated with processor management means that shared-memory systems do not scale to a large number of processors, making it a less attractive option for large-scale parallel computing than distributed memory systems. In this case, processors only have access to their local memory space, and communicate with each other via messaging passing. Information to be sent and

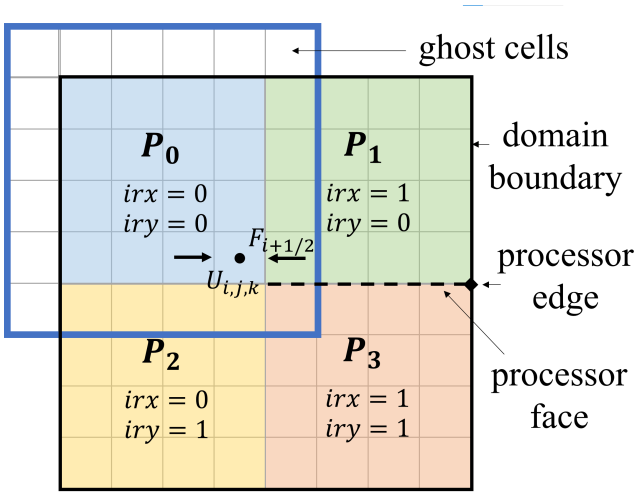


Figure 3.1.: Demonstration of domain decomposition. Full spatial domain (black outline) is split into four sub-regions which are sent to four corresponding processors $P_0 - P_3$. Shaded blue cells indicates the cells on which the solver is applied for P_0, blue box around P_0 shows the surrounding ghost cells that are also required by the solver. These overlap with cells in P_1, and therefore require inter-processor communication.

received are copied into temporary stores, then sent across processors via an interconnection network, rather than being directly accessed. While this approach is attractive for its scalability, communications via this architecture have a longer associated latency than the shared-memory case. A code's performance will therefore depend both on algorithmic design and physical architecture.

MPI library

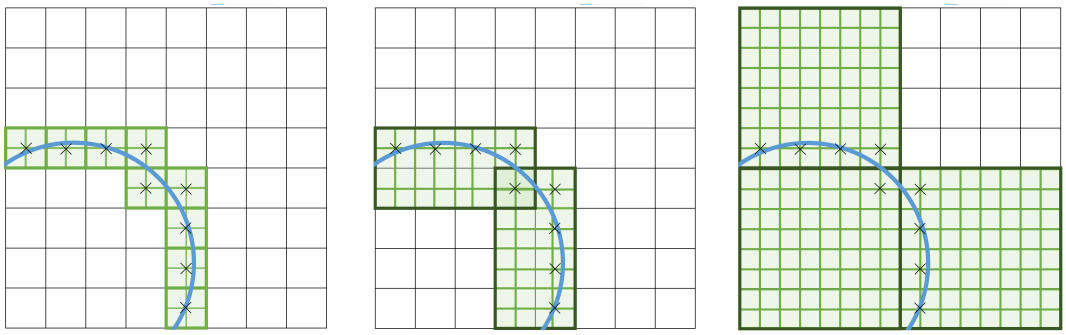
The framework most commonly employed to communicate data across distributed memory systems is the Message Passing Interface (MPI) library [78]. Each processor is initialised with a unique identifier called a *rank* in the global space, but can also have identifiers within sub-groups of processors called *communicators*. Virtual *topologies* to define processors in relation to each other can be generated, for example Fig 3.1 conforms naturally to the 'MPI Cartesian' topology provided by library. Communicators can be formed between processors in the same rows, $[P_0,\ P_1]$ and $[P_2,\ P_3]$, where P_0 and P_2 would have the identifier $irx = 0$ within their respective communicators. The routines provided by the MPI library enable efficient communications either directly between two processors (point-to-point communications), or simultaneously across several processors (collective communication). Communications can also either be blocking, where the processors cannot proceed with the rest of the program until the communication has been successfully executed, or non-blocking, allowing processors to complete other tasks and checking the communication status later. Examples of these are provided in §3.4.

Ghost cells

§2.1.1 describes how the set of hydrodynamic and MHD equations are solved on a discretised grid using finite-differencing schemes. Eqs 2.17 and 2.21 show that for each cell $U_{i,j,k}$ information from its neighbouring cells are required to calculate the fluxes across its six faces. Using simple upwinding (Eq 2.21) requires only a single neighbour on the right hand side to calculate $F_{i+1/2,j,k}$, whereas using a higher order correction (Eq **??**) requires two neighbours. Consider the rightmost column of cells of P_0 in Fig 3.1. These require one or more adjacent columns of cells from P_1 to construct the leftward fluxes on the right faces. The full domain of dependence for P_0 for simple upwinding is therefore shown in the blue outlined box, which constitutes the domain cells where calculations are performed (shaded blue cells), and their surrounding 'ghost' or 'halo' cells which are populated with information received from neighbouring processors, i.e. $P_1 - P_3$. Likewise, P_1 will have its own ghost cells populated through communications with processors P_0, P_1 and P_3. As alluded to above, the latency associated with MPI communications makes it desirable to reduce the number of calls, therefore the 'corner' communications (P_0 communications with P_3 and vice versa) can be avoided as described in §3.4.

3.3. Review of refinement strategies

The three primary refinement strategies are described here with respect to Fig 3.2. This shows a uniform mesh overlaid with a curved feature that might represent, for example, a material interface from a supernova explosion such as in Fig 1.4, which requires high resolution to capture the formation of fine-scale instabilities. In adaptive codes, 'refinement indicators' such as gradients of density or velocity are used to identify these features as they evolve. Cells that have an indicator value above a specified threshold are tagged for refinement, such as in Fig 3.2 represented by black crosses. The resultant grid topology after refining these cells at one additional level of refinement using the three strategies are then shown, where shaded green regions indicate higher resolution cells and dark green outlines indicate patch or block boundaries.



(a) Cell-based splitting (b) Patch-based splitting (c) Block-based splitting

Figure 3.2.: Illustration of three refinement strategies with respect to a curved feature.

3.3.1. Cell-based refinement

As the name suggests, cell-based refinement splits the grid on a cell-by-cell basis, making it the most targeted form of refinement. This is shown in Fig 3.2(a), which visibly produces the least number of high resolution cells of the three strategies. This makes it the most flexible method for providing resolution along complex boundaries, while minimising the number of grid computations.

The resultant topology naturally lends itself to an 'octree' data structure; a hierarchical structure where each node (cell) has a 'parent' (coarser) cell and two/four/eight 'children' (more refined) cells in 1D/2D/3D, illustrated in Fig 3.3. This 2D toy example shows a base grid of 2×2 cells with the upper-left corner refined twice, providing four times the base resolution. The corresponding tree representation is shown on the right. In this example, the green cell with index 2 has a coarser parent (blue cell index 2) and four children (yellow cells). Any node without children are considered 'leaf' nodes (outlined in black in tree diagram), and form the simulation grid. Tree structures are therefore useful for storing hierarchies where children nodes are generated from exact multiples of parents, however there can be significant memory and time overheads when using them for cell data storage. Uniformly spaced grid variables are typically stored in

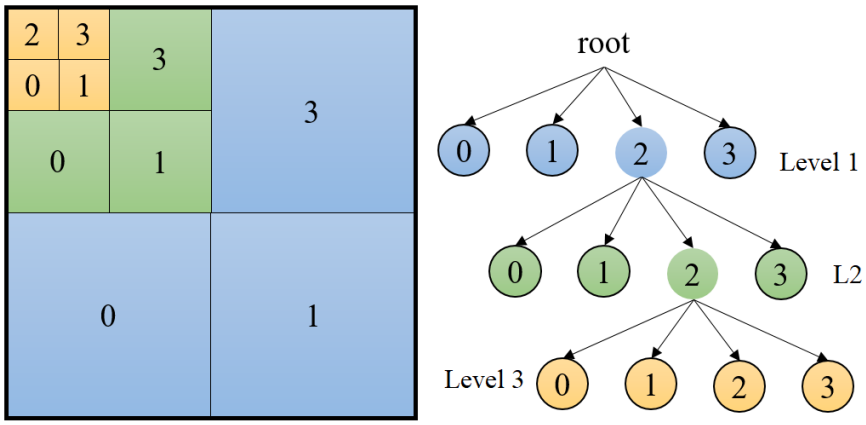


Figure 3.3.: Demonstration of cell-based splitting the top-left corner of a 2×2 cell grid, to reach four times the original resolution. The corresponding tree structure is shown on the right, where leaf nodes are circled in black and form the simulation grid.

arrays, which are linear data structures stored contiguously in memory space. A logically rectilinear grid which uses the indexing notation i, j, k for the three spatial dimensions therefore allows for straightforward accessing of neighbouring cells $(i \pm 1, j \pm 1, k \pm 1)$. This is lost when using tree structures since cells do not necessarily have a single neighbour per face.

Two options then remain to access neighbours. The first is performing tree-traversals, i.e. navigating through parent/children networks until the desired cell is reached. The binary search tree is an example of a balanced tree that minimises search times, however this still scales as $\mathcal{O}(N)$, where N is the number of nodes. Performing a time-consuming search operation every time neighbouring cell data is required is clearly not ideal. Recent improvements in hashing techniques appear promising for faster access as shown in Ref [80], however to the author's knowledge this has not yet been implemented in prominent hydrodynamic or MHD codes. The second solution is storing memory addresses of neighbours once the tree structure is formed. This requires 6 additional bytes per cell for each immediate neighbour, potentially more if the reconstruction scheme is higher order and requires a larger stencil, resulting in a large memory overhead. Solutions such as the FTT (fast threaded tree) structure developed by Khoklov can provide memory savings by grouping cells into 'octs', and is used in several cell-based refinement libraries [81] [82].

One such example of a cell-based AMR code is RAMSES, an open source cosmological code used to simulate "self-gravitating, magnetised, compressible, radiative fluid flows" [83]. This uses the FTT structure in an adaptive implementation for the hydrodynamic solver, and is coupled to a particle solver. It has been used for simulations utilising 256^3 particles, where a uniform grid to the same highest resolution, or "equivalent resolution" grid, would require 8192^3 particles.

3.3.2. Patch-based refinement

This method involves constructing overlapping 'patches' of high resolution around cells that have been marked for refinement, as demonstrated in Fig 3.2(b). This was the technique used in the first documented form of mesh refinement [58], and is still used by several codes today. A patching algorithm is employed which uses edge detection and clustering techniques to construct rectangular patches over cells tagged for refinement.

The key advantage to this method over cell-based refinement is maintaining the ability to perform computations on logically rectilinear grids. An array can be constructed per patch, which is operated on in the same way as the uniform grid. This requires less invasive changes when upgrading a code to include refinement, and retains fast neighbour access. The flexibility of defining patches of arbitrary size also makes it amenable to refining irregular regions. A visible drawback, however, is a larger number of cells and less targeted resolution than cell-based refinement, even where it is not required. As such, this method is better suited to sparse regions of high resolution. Furthermore, because grids of different levels can overlap the solution needs to be consistent across resolutions. The communication across levels can be an additional challenge, as higher level patches may not have a unique parent patch. Finally, in adaptive codes the clustering algorithm is frequently performed, leading to a large overhead that can offset the variable resolution savings.

Despite the above challenges, significant development work has been done by the Exascale Computing Project (LBNL, NRL, ANL) to develop an open source, exascale, patch-based mesh refinement library called AMReX [84] [85]. It has been integrated into several astrophysical solvers, including Castro and Maestro, and has demonstrated impressive scaling for hydrodynamic solvers on massively parallel systems of order $\mathcal{O}(10^5)$ processors.

3.3.3. Block-based refinement

The block-based approach combines the recursive splitting of cell-based refinement and the rectilinear nature of patch-based refinement. In this case, the domain itself (rather than an individual cell) is recursively split until a combination of low and high resolution 'blocks' or grids cover the domain. Unlike patch-based refinement, the blocks are not overlapping, avoiding the need for different solutions at coincident locations. A tree structure as described previously can be used to organise the connectivity of blocks, which will have dimensions of the order of the number of processors, rather than number of cells. Access to neighbouring processors nodes is also only required during halo exchange (discussed in §3.4.1), rather than when constructing finite-difference stencils on a cell-by-cell basis. Despite the simplicity of this approach it also suffers from over-refinement of the domain, as shown in the extreme example in Fig 3.2(c). It significantly underperforms relative to cell-by-cell refinement and efficient clustering algorithms. However, in the limit of many refinement levels, the fraction of over-refined cells will diminish.

This scheme has been used in several other codes, including FLASH (Paramesh library), ENZO, and Athena++ [86], [87], [88]. The complexity of mesh refinement means details of the implementation differs between codes, however the fundamental concept is the same across them. The implementation in Chimera will be outlined in more detail in subsequent sections.

3.3.4. Load balancing

The benefits of parallelisation are reduced if processors are idle while waiting on another processor before they can complete their own tasks. Load balancing refers to ensuring that each processor performs a roughly equal number of computations, such that no single processor is a bottleneck. For a uniformly-spaced grid this is easily achieved by allocating an equal number of cells to each processor, although for non-uniform grids this is trickier. A commonly used method for tree-based data structures is to generate a space-filling curve (SFC), which provides a mapping of unstructered data to a linear array, for example the Z-order Morton curve [89] shown in Fig 3.4(b). The SFC can then be equally split and distributed across processors.

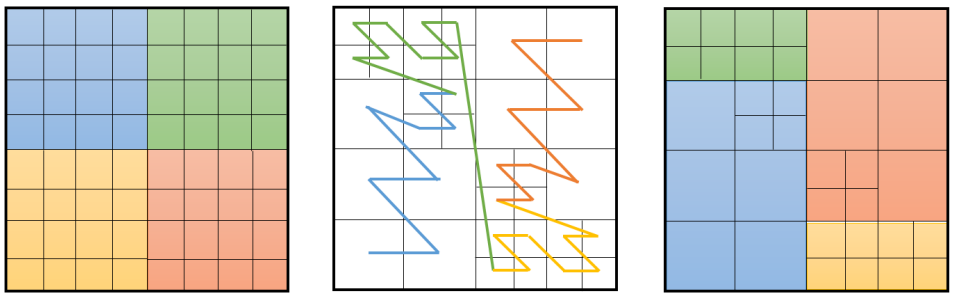


Figure 3.4.: Demonstration of load balancing across processors for a uniform mesh (left), and an unstructured mesh mapped by a Z-order Morton SFC (right).

3.4. Previous implementations

Successive improvements to the mesh refinement capability were built on concepts from previous versions, therefore a chronological overview of the development is provided. Standard parallelisation of a uniform mesh is first described, followed the first implementation of a two-level grid by previous developers, and finally an extension to the two-level version, which ultimately provides context for the final implementation described in §3.5.

3.4.1. Uniform mesh

In the absence of mesh refinement, Chimera uses a logically rectilinear grid as described in §2.1.1. To recap, the full spatial domain is split into $N = N_x \times N_y \times N_z$ cells, which can each be uniquely identified by the set of indices i, j, k for each dimension. For a large value of N, the grid can be parallelised into $nrx \times nry \times nrz = n_{proc}$ processors using an MPI Cartesian connectivity topology. Processors are grouped together via communicators into $\hat{x}$, $\hat{y}$, and $\hat{z}$-direction groupings. Each processor can therefore identified in terms of its *rank* in the global communicator, and the *irx*, *iry*, and *irz* indices for each of the three communicators, see Fig 3.5. The number of cells per processor is then $n = n_x \times n_y \times n_z$, where $n_x = N_x/nrx$.

To populate the ghost cells in each processor, Chimera uses four sets of point-to-point blocking communications in each direction. Specifically, these are pairs of `mpi_send` and `mpi_recv` calls. Note, assuming non-periodic boundary conditions processors that lie on the domain boundary will not send or receive data in the direction of the boundary. In the $\hat{x}$-direction the four steps are as follows:

1. Even *irx* sends to *irx+1* neighbour, odd *irx* receives from *irx-1* neighbour,
2. Odd *irx* sends to *irx+1* neighbour, even *irx* receives from *irx-1* neighbour,
3. Even *irx* sends to *irx-1* neighbour, odd *irx* receives from *irx+1* neighbour,
4. Odd *irx* sends to *irx-1* neighbour, even *irx* receives from *irx+1* neighbour,

The staggering of send/receieve calls in terms of processors with even or odd values of *irx* is required to prevent *deadlock*. This occurs for blocking communications when a program cannot proceed because processor A is waiting for a communication with processor B to be completed, which itself is waiting for another communication.

As seen in Fig 3.1, the ghost cell region around P_0 also includes corner cells from P_2. Rather than implement additional sets of corner MPI transfers, the number of cells exchanged in the $\hat{x}$-direction (e.g. between P_0 and P_2) can be extended from $n_y \times n_z$ cells to $(n_y + 2) \times (n_z + 2)$ cells. Likewise, transfers in the $\hat{y}$-direction can be extended to $(n_x + 2) \times (n_z + 2)$ cells. Performing these extended sweeps in all three directions automatically ensure that all corner cells are exchanged. Note, this is a straightforward procedure for uniform grids but not for variable mesh grids.

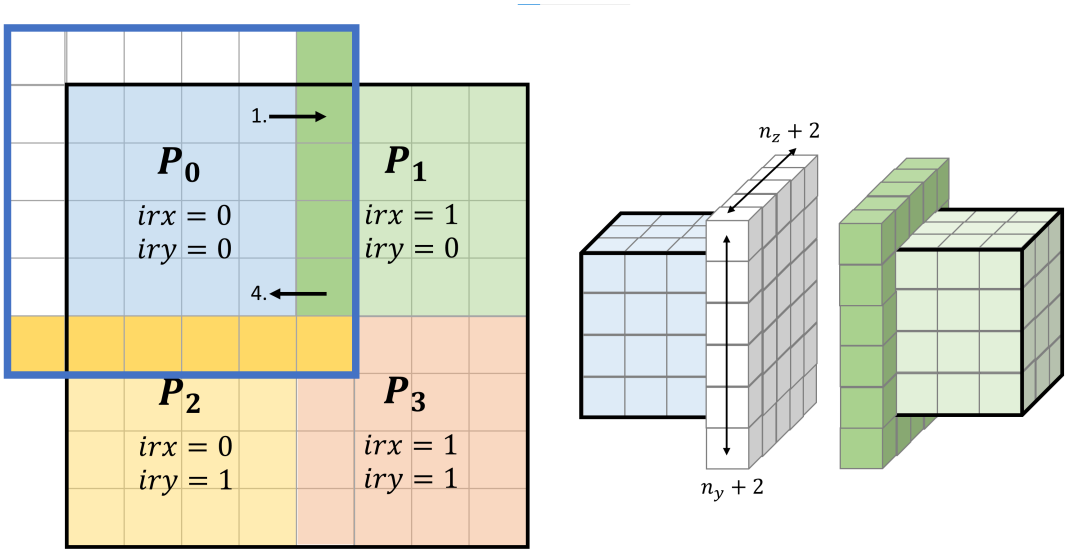


Figure 3.5.: Demonstration of communications between processors in the $\hat{x}$-direction to populate ghost cells on $\hat{x}$-face. *Left:* 2D representation, where processor numbers are labelled in large font, underneath are respective *irx* and *iry* indices assuming a Cartesian topology. Arrows correspond to steps 1 and 4 in the series of point-to-point communication. *Right:* 3D representation, showing transfer of information from P_1 to P_0 using an extended region. Performing steps 1-4 described above for the x and y direction with this extended region ensures that the bottom right corner of P_0 is populated with information from P_4, without requiring an explicit transfer.

3.4.2. 2-level method

The first version of a variable resolution grid in Chimera was conceptualised and implemented in 2013 by C. Jennings and J. Pecover to study ETI growth in MagLIF liners [10]. This version used a single level of refinement in one region along the $\hat{x}$ direction, and is best understood using Fig 3.6 from Ref [10].

In this demonstration, the full domain is split into three processors along the x direction, labelled P_0, P_1, and P_2, each with $n = n_x \times n_y \times n_z$ cells and cell volumes $\Delta V = \Delta x \times \Delta y \times \Delta z$. This forms the low-resolution grid. Consider the case where the region covered by the parent processor P_1 is resolved to twice its resolution. Two children processors are created in each dimension, resulting in eight additional processors P_3-P_{10}, which form the high-resolution grid. Each child processor has the same number of cells as its parent, but with half the spatial extent in each dimension and therefore one eighth the cell volume, $\Delta V_{\text{child}} = \Delta x/2 \times \Delta y/2 \times \Delta z/2 = \Delta V/8$. This is an example of block-based refinement, as the splitting of P_1 generates further blocks of cells with the same dimensions. To define the refinement region, lower resolution processors are tagged as 'entry' processors (first to be refined) and 'exit' processors (next one unrefined), in this example P_1 and P_2 respectively.

The two resolution grids therefore evolve almost entirely independently, undergoing intra-grid communications as described for a single resolution case above. After completing all internal communications, a bespoke set of transfers is then executed for processors adjacent to neighbours of different resolutions. These are discussed in further detail later in the context of improvements to this implementation.

This method has two significant advantages:

1. Load balancing is achieved by default as all processors have the same number of cells and therefore execute the same number of operations in the core solver.

2. Modifications to the existing code are minimal. The physics subroutines typically loop around $n_x \times n_y \times n_z$ cells, which is unchanged for all processors in this method. The only additions are (i) defining where the entry and exit processors are, (ii) implementing inter-grid exchange routines, and (iii) outputting an additional grid.

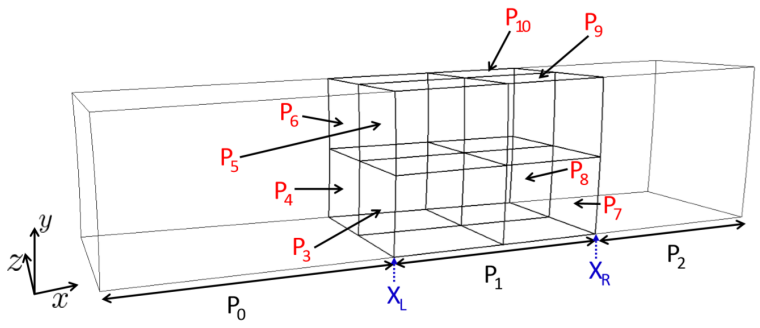
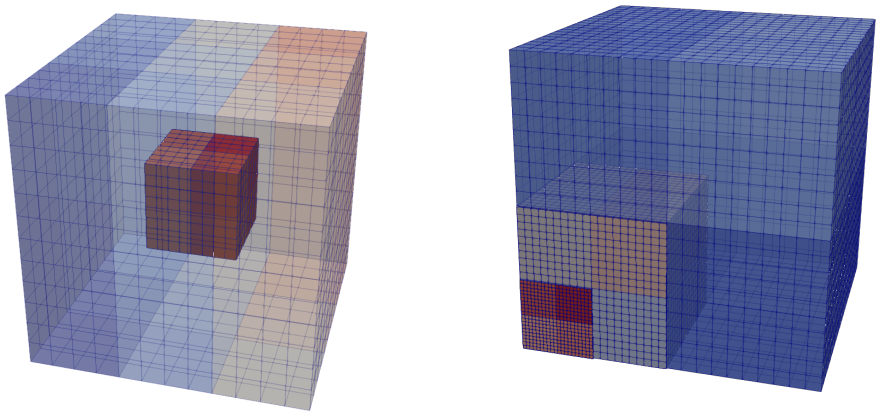


Figure 3.6.: Demonstration of initial implementation of single-level mesh refinement in Chimera. Reproduced from [10].

3.4.3. N-level method

A single level of refinement and splitting in a single direction, however, is too limited for most other problems of interest. As such, the natural next step was to extend this method to the $\hat{y}$ and $\hat{z}$ directions, and to higher number of refinement levels (Fig 3.7). For l refinement levels, the factor of maximum resolution is 2^{l-1}. Although the core concept remains the same, this required a complete rewrite as existing routines were hard coded for two refinement levels.



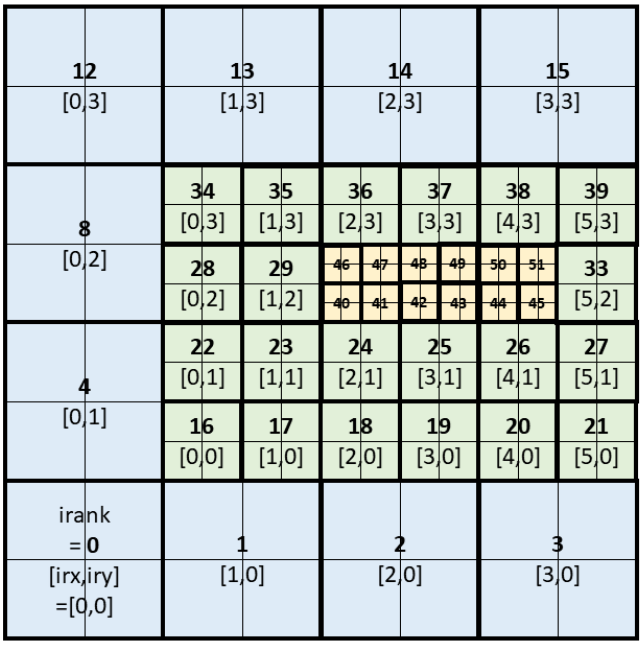
(a) (b)

Figure 3.7.: Extension of implementation in Ref [10] to 3D and arbitrary number of levels, demonstrated in two refinement configurations; (a) an extension of Fig 3.6 to 3D (b) example of three grid levels.

As mentioned above, the modifications for integrating this capability are i) initialisation of processor topology, ii) ghost cell exchange with neighbouring processors, and iii) output of higher resolution grids. These are described here in this order.

- **Initialisation**

  The single-level, single direction version identified 'entry' and 'exit' processors on the low resolution grid to define the extent of the high-resolution grid. A similar concept was employed in this extended version, except for all three dimensions, and with respect to different grid levels. This is best demonstrated using Fig 3.8, which shows three levels of refinement in 2D for simplicity. For a given grid (identified by $igrid$), the higher resolution grid ($igrid + 1$) is defined in terms of x and y entry and exit processors. The number of higher resolution processors in each direction is then given by this difference times 2, e.g. nrx for $igrid = 2$ is given by $(x_{\text{exit}}^{igrid=1} - x_{\text{entry}}^{igrid=1} - 1) \times 2 = 6$. An evident drawback of this implementation is the complexity associated with defining grids at increasingly high resolution, for example initialisation of a 6-level grid requires knowledge of the grids at all levels below. Note, neighbouring processors are only permitted a difference of one resolution level.



igrid	1	2	3
nrx	4	6	6
nry	4	4	2
x entry	1	2	-
x exit	4	5	-
y entry	1	2	-
y exit	3	3	-

Figure 3.8.: Example splitting of domain into three levels of refinement, using the extended implementation of mesh refinement. *Left:* processors (thicker black outline) with 2×2 cells shown with $irank$ and irx/iry indices labelled. Cell shading shows grid level and therefore resolution level. The blue grid has lowest resolution (largest cells), and the yellow grid has the highest resolution (smallest cells). *Right:* Table describing properties of each grid, including number of processors per direction and entry/exit points for higher level grid. Note, underlying processors in $igrid = 1$ and $igrid = 2$ are not shown for clarity.

Using this configuration each grid is initialised as a separate Cartesian topology with its internal communicator, such that the $irx/iry/irz$ indexing applies within each grid. In the global domain, each processor still has an unique *rank* identifier, which is used to store neighbour information across grid levels. For example, processor P_4 in Fig 3.8 will have 'child' neighbours P_{16} and P_{22}, which themselves will store P_4 as a 'parent' neighbour.

- **Ghost cell exchange**

  In §3.4.2, communications between the two grids were performed using point-to-point (direct `mpi_send` and `mpi_recv`) calls as only a small number of high-resolution processors were created in this process. In the generalised case which produces a much larger number, a set of more efficient collective operations to distribute and gather the data were used instead. Fig 3.9 shows this type of transfer in 3D at a resolution interface. The low-resolution parent neighbour packs its data into a linear array (as required for MPI routines) and uses the `mpi_scatterv` collective routine to distribute the data to the four children neighbours. The reverse

process is the `mpi_gatherv` routine, which collates data from different children processors and sends it to the parent processor. Details of the interpolation routines for populating the send/receive data are described in §3.5.2.

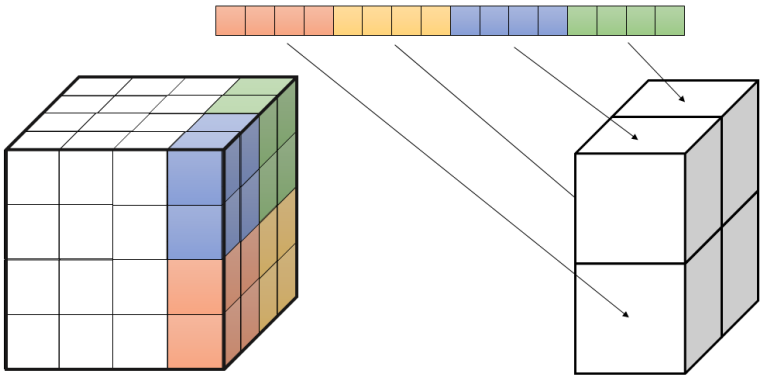


Figure 3.9.: Distribution of data from a low resolution 'parent' processor on the left, to four higher resolution 'children' processors on the right. Communication performed using the collective `mpi_scatterv` routine.

- **Output** Because the underlying parent processors are retained in this method, each grid can be separately written out as fully populated uniformly spaced array. When visualising data from the full simulation, care must be taken to ensure the correct solution from the higher resolution grids are being viewed rather than the incorrect solution from the underlying, low resolution processors.

This version of the code was found to work successfully on integrated problems of interest. For example, a 3D simulation of the final stagnation phase of a perturbed capsule implosion (Fig 3.10 was used to test the behaviour of the code. The simulation was initialised from a uniform grid case just prior to stagnation, and allowed to evolve into the re-expansion phase. Two levels of refinement were used, as shown in the left images. The plots on the right are lineouts through the $\hat{x}$ axis on each grid, where the orange, green, and black lines correspond to grids of increasing resolution. The density profiles (top right) show the shell reaching a more compressed state at higher resolutions. The temperature profiles (bottom right) also show the fuel reaching higher peak temperatures with increasing resolution. The lineouts on the right show that the transitions between different resolution grids is smooth, giving confidence in this method. In the later re-expansion phase, the signals move from higher to lower resolution, therefore the inter-grid transitions remain smooth however the lowest resolution lineout (orange line) becomes discontinuous as information from the underlying processors are overwritten.

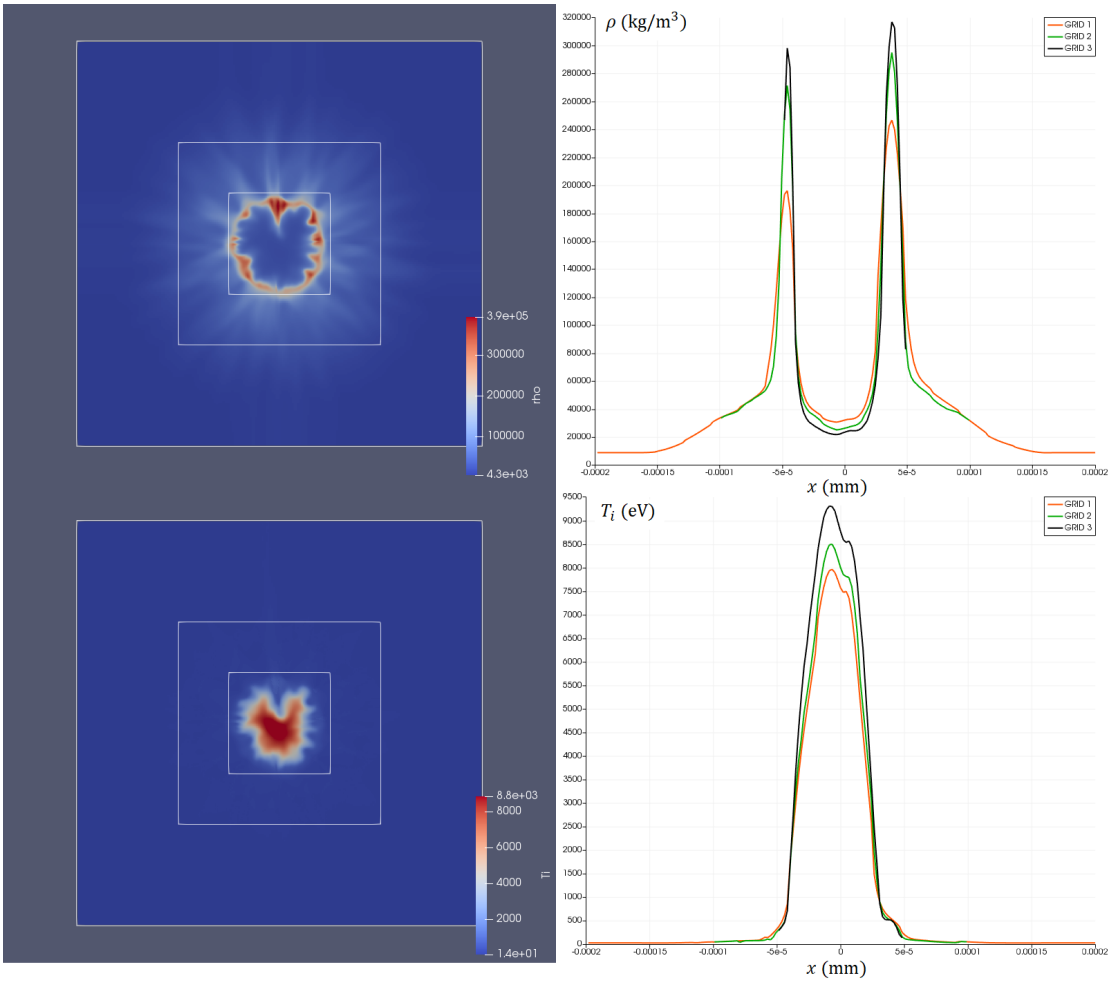


Figure 3.10.: *Left:* Slices through a 3D simulation of capsule implosion with an applied multi-mode perturbation using three levels of mesh refinement, where the domain boundary of each grid is outlined in white. *Right:* lineouts through the $\hat{x}-$axis for the three grid levels (orange, green, and black line for $igrid = 1, 2, 3$ respectively).

Final remarks

While this implementation of the refinement was demonstrated to work in cases such as the above, it was still found to be quite inflexible, even with the aforementioned extensions. The design requires processors to be split in all three dimensions, forcing a large number of processors to be generated even for 1D or 2D problems. Furthermore a lot of redundant computation is performed by retaining the underlying parent processors which do not contribute to the final solution. Additionally, the design allowed only for a single region of refinement, i.e. only one grid per level, which is limits its application. As such, another re-design was carried out, retaining the core concepts but introducing large-scale improvements.

3.5. Current implementation

The existing implementation in Chimera includes the benefits of the above method, namely load balancing and minimal modifications to physics routines, but addresses its issues; redundant computation in overlapping processors, and flexibility to refine in 1/2/3D. A library called *stare* (static refinement) was developed to provide the following features as part of a mesh refinement capability:

(i) grid splitting, in a more flexible and user-friendly manner

(ii) ghost cell exchange, across same and different resolution-level processors

(iii) data output format (I/O), for easier visualisation of data.

These are described in this order in the following sections, with the exception of ghost cell exchange for MHD variables, which requires a more rigorous treatment and is covered in (§3.6).

3.5.1. Grid splitting

The data structures and block-based refinement algorithms used to generate the grid topology are described here. A tree-style hierarchy, where each node is a processor or 'block' with an equal number of cells, is generated by recursively splitting blocks until a refinement criteria is met. The refinement criteria is specified using a derived type called 'refinement shape', with which users can define certain geometries and an associated resolution. The leaf nodes of the tree then form the grid structure on which the solver works. The splitting algorithm relies on two main data types:

- `block_type`: stores processor information and can generate children blocks
  - Contains spatial information, including its origin, number of cells, cell resolution. From this its spatial extent (x_0, x_1, etc.) can be defined.
  - Stores a pointer to the 'root' block, and pointers to children blocks. A block is defined as a leaf node if its children are unallocated.
- `refinement_shape_type`: provides a way to define geometric refinement criteria
  - Stores topological information. For example, if the refinement shape is a cuboid, then its upper and lower extents in each direction are specified.
  - Also has a minimum resolution to which the shape must be resolved. This allows several shapes to be refined at different minimum resolutions.

Using these two structures the algorithm then works as follows (illustrated also in Fig 3.11).

1. *Initialisation:* The user defines a a base grid (block type) which spans the full simulation domain and has the lowest resolution cells. The base grid is itself split into $n_x \times n_y \times n_z$ blocks, which dictates how many cells are contained in each processor. Additionally, the user generates a list of refinement shapes, with their respective minimum resolutions.

2. *Recursive split:* The base grid and list of shapes are passed to a recursive splitting algorithm, which generates a full tree hierarchy. At each block, all refinement shapes are checked for two criteria; if they intersect the block, and if the block's cell spacing is less than or equal to the shape's minimum resolution. If these are true and false respectively, then the block is split.

3. *Check consistency:* Once the initial tree structure is formed, the blocks are checked to ensure their neighbours are different by no more than one resolution level, otherwise further refinement occurs. This typically takes a few iterations through the full structure to ensure the grading process is converged.

4. *Assign and delete:* The number of leaf blocks from the converged structure is checked to match the number of MPI processors initialised. If so, each processor is assigned a block from which its spatial information and neighbours are assigned. After this point the tree structure is no longer required, and is deleted.

Block splitting

Using the above data type, a block is split by allocating an array of children which are of the same derived type [2]. The dimension of the children array depends on whether the user defines the grid to be refined in $\hat{x}$, $\hat{y}$, and $\hat{z}$. The example in Fig 3.11 shows refinement only in two directions, therefore each parent only produces four children. For each dimension that is split, the children's resolution in that direction will be halved and their origin appropriately allocated. Children blocks also inherit information about whether they lie on the simulation domain boundary, for example in Fig 3.11, blocks 2 and 3 in green will be flagged as a boundary processor in the upper $\hat{x}$ direction.

Tree-traversals

For a given tree structure, different logic can be used to determine the order in which nodes are accessed. The initial recursive splitting of blocks is best suited to a pre-order traversal, which moves as far down the left sub-tree as possible before moving back upwards and to the right (Fig 3.12(a)). Therefore if a block is split, its children are recursively split before a neighbouring block of the same level is considered. However using this logic when deleting the structure would result in deallocating parent blocks and rendering pointers to children blocks inaccessible, therefore a post-order traversal must be used which moves from the leaf nodes upwards to the root (Fig 3.12(b)).

[2] Requires Fortran 2003 and above.

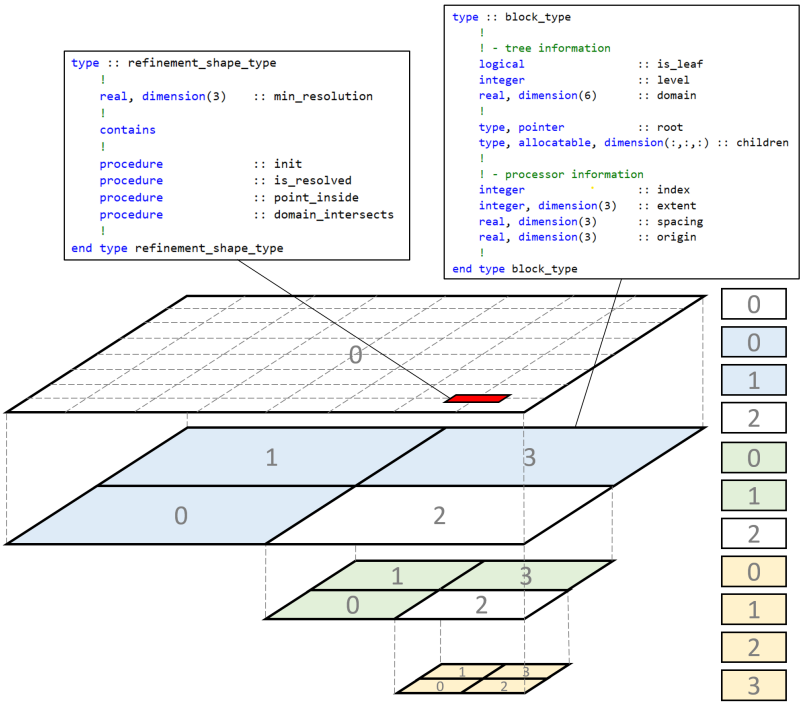


Figure 3.11.: Illustration of the grid splitting procedure. Red box on the base grid indicates region of interest to be refined. Remaining tree hierarchy of blocks shown underneath, where shading of leaf blocks corresponds to resolution level. Traversal order to split the grid shown on the right. Data structures used to define refinement shape (red box) and blocks shown in insets.

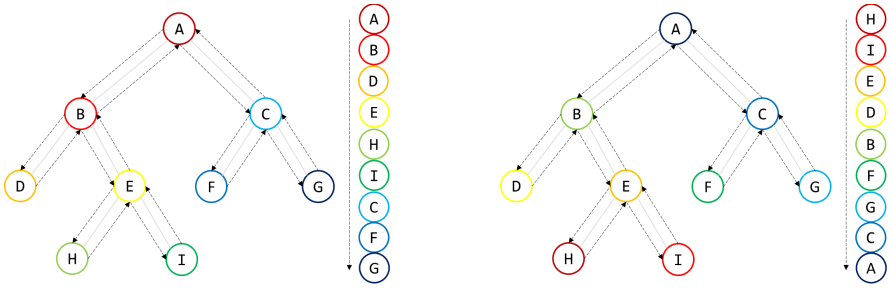
Intersection logic

One possible refinement shape is a cuboid, as shown in Fig 3.11 for a 2D simulation. The logic to determine if a cuboid intersects a block (another cuboid) is trivial. Another useful refinement shape is a cylindrical shell, which requires more complicated logic to test for a block intersection, as shown in Fig 3.13. Usage of both is demonstrated in Fig 3.17 for pulsed-power driven magnetic reconnection (see §6), which uses a cuboid for the central reconnection layer and cylindrical shells for the two wire arrays.

Final remarks

For complicated topologies the above procedure is typically run on a single processor first to generate the number of leaf nodes, and re-run on a multi-processor run with the specified number of cores. The CPU time taken to generate, for example, a tree with 3064 blocks and 8 levels of refinement is < 0.14 seconds on an AMD EPYC 7742 processor. As this is only performed once on initialisation, the time taken to construct the grid is negligible on the scale of simulation run time.

Three further avenues of improvement remain for this tool. Firstly, while it is readily extended to cylindrical $\hat{r} - \hat{z}$ geometry as demonstrated in §5.3, extensions could be developed to make the tool geometry agnostic, i.e. for $\hat{r} - \hat{\theta}$ geometry or spherical



(a) Pre-order traversal (b) Post-order traversal

Figure 3.12.: Example traversal orders for navigating a tree structure. Both orders show the tree structure on the left, and the order of the nodes accessed using their respective logic on the right.

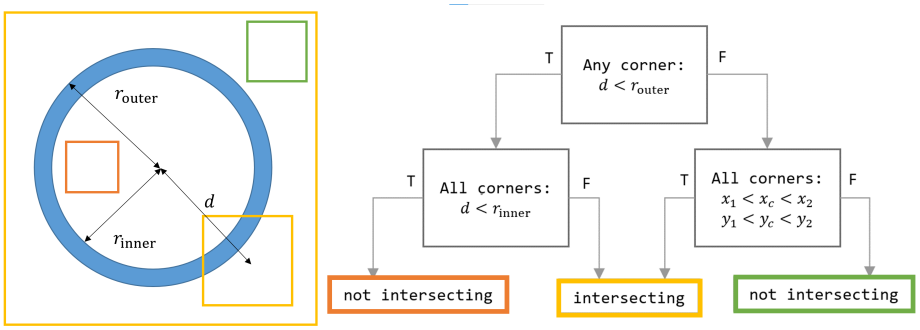


Figure 3.13.: *Left:* example cylindrical shell refinement object, with possible configurations of intersecting and non-intersecting blocks. *Right:* logic used to determine whether a block and cylinder intersect, given their dimensions.

geometry, where grids are resolution-starved in the angular directions at large radii. Secondly, only cuboids and cylindrical shell refinement shapes are currently available, as these are the most commonly used geometries for typical pulsed-power applications. Development is underway to include refinement from STL (stereo-lithography) files, which define closed surfaces for arbitrary geometries using a series of triangular facets. This can be used to include large scale structures in pulsed-power systems such as electrodes or return-posts.

Thirdly, there are potential ways to extend the splitting to be adaptive rather than static. In an adaptive code, the refinement at a later time in the simulation will likely result in a different number of leaf nodes than at initialisation, which poses a challenge since the simulation job is submitted with a fixed number of processors. This may force the user to each submit a job with an excessively large number of processors as a conservative estimate for number of nodes produced, or require a single processor to solve two or more blocks. The latter option would forego the natural load balancing of a static capability where each processor is assigned exactly the same number of cells. Another consideration is the criteria used for refinement; these are currently objects (shapes) defined independently of the grid structure. In an adaptive code a set of refinement indicators would need to be specified from the grid variables (discussed in §3.3), adding further complexity.

3.5.2. Ghost cell exchange

MPI communications

As part of the mesh refinement capability, an upgraded set of MPI routines were also provided. The use of collective operations as described in §3.4 provides performance improvements when, for example, a parent processor sends data to multiple children processors simultaneously. However these were not used in the current implementation, as the directionally-flexible splitting gives one, two, or four neighbours making the collective routines more challenging to implement.

To improve performance non-blocking point-to-point communications were implemented instead. Non-blocking (NB) routines allow a communication to be initialised but only checked for completion at a later time, therefore data can be requested by processor A from processor B, but A can be allowed to progress with other tasks until the data is explicitly required. This removes the need for left/right staggering of send/receive communications to avoid deadlock, as is required for blocking communication (§3.4.1). Each NB send and receive command (`mpi_isend` / `mpi_irecv`) is paired with a corresponding `mpi_wait` command to ensure the data has been received, and are most optimally used in the following order:

1. call `mpi_irecv` from left and right neighbours
2. Pack data, and call `mpi_isend` to left and right neighbours
3. Check status of receive (`mpi_wait`). If completed, unpack data
4. Check status of send (`mpi_wait`). If completed, exit.

By calling the NB-receive at the start, a processor may receive data in its temporary storage while completing step 2, after which it can immediately progress to unpacking the received data. A test program was performed on 128 AMD EPYC 7742 processors on the ARCHER2 supercomputing facility to exchange 100 bytes of data to left/right neighbours 1000 times. This took an average of 95 milliseconds using blocking communications and 42 milliseconds using NB ones, showing a substantial performance improvement.

Hydrodynamic variables

This section describes how to populate a processor's ghost cells with cells from neighbouring processors, for cell-centred variables. Three different scenarios can arise, illustrated in Fig 3.14:

1. neighbouring cells are of the same resolution,
2. neighbouring cells are of lower resolution,
3. or neighbouring cells are of higher resolution.

In the first case the layer of neighbouring cells can be simply injected directly into the ghost cells, which is effectively equivalent to all cells being solved on the same processor. The second case is referred to as 'prolongation', and requires an interpolation of the parent cell values to the ghost cell locations. The third case is referred to as 'restriction', and requires averaging of the finer (higher resolution) cells to the ghost cell locations. The latter two terms derive from tree-based hierarchies for AMR codes, where additional refinement requires extending or prolonging a tree structure, and coarsening cells requires reducing or restricting the tree.

In the latter two cases, two conditions must be satisfied to ensure conservation of U and the flux of U across an interface [58] [87]:

$$U^{\text{coarse}} = r^{-d} \sum U^{\text{fine}}, \tag{3.1}$$

$$F^{\text{coarse}}_{i\pm 1/2,j,k} = \sum F^{\text{fine}}_{i\mp 1/2,j,k}, \tag{3.2}$$

where r is the refinement ratio (in this work $r = 2$ throughout) and d is the dimension of the simulation. In the case of prolongation the first criteria can be achieved using a simple zero-order interpolation, i.e. $U^{\text{fine}}_{\text{ghost}} = U^{\text{coarse}}_{\text{neighbour}}$. For restriction in three dimensions, an eight-cell average of the finer cells is required to populate a single coarse cell. This, however, does not necessarily guarantee that the sum of fluxes exiting finer cells is equal to the flux entering a coarser cell, therefore the additional correction in Eq 3.2 must be applied. This scheme was implemented in the mesh refinement code without a flux correction step and seen to behave in a stable manner without it. As such, simulations presented are run without this correction, however it will be implemented in future to ensure complete accuracy.

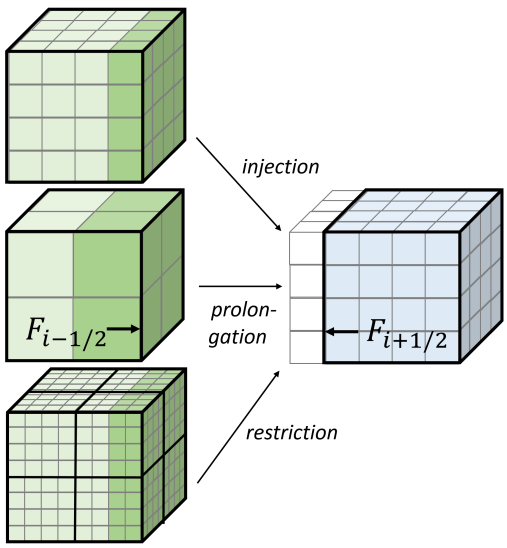


Figure 3.14.: Demonstration of injection, prolongation, and restriction operations to populate ghost cells in the processor on the right (white cubes). Extended extents demonstrated in Fig 3.5 not shown here for clarity.

Interpolation

With the flexibility of directional splitting, a given cell can have one of the seven split configurations shown in Fig 3.15, which the prolongation and restriction routines for sending data across different resolution processors must account for.

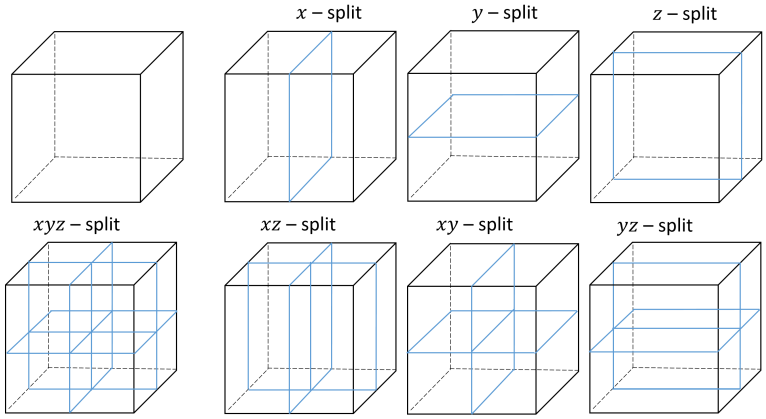


Figure 3.15.: Seven possible splitting configurations using a flexible splitting approach

Fig 3.16 additionally shows how low resolution cells are indexed in order to populate data for the high resolution cells. In order to go between indices ix/iy/iz (fine cells) and ix1/iy1/iz1 (coarse cells) for a generalised case, the following generalised expressions can be used:

$$\text{ix1} = \frac{\text{ix1} + x_m - 1}{x_m} \qquad \text{iy1} = \frac{\text{iy1} + y_m - 1}{y_m} \qquad \text{iz1} = \frac{\text{iz1} + z_m - 1}{z_m}, \tag{3.3}$$

where fractions indicate integer division. While this indexing requires multiple arithmetic operations to access a single cell, this is performed in favour of explicitly writing prolongation/restriction routines for all seven configurations.

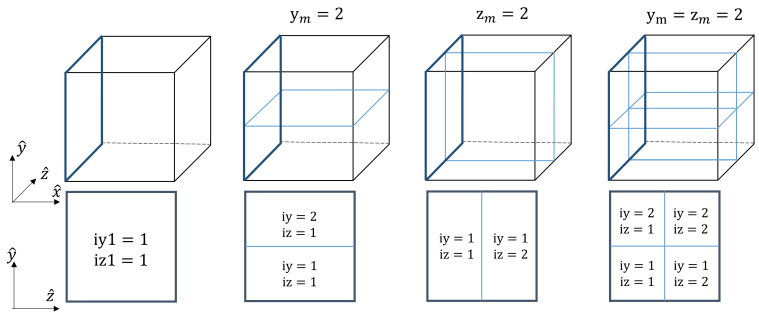


Figure 3.16.: Interpolation indexing for different splitting configurations, shown for $\hat{x}$-direction ghost cells where x-indexing is not relevant. Leftmost cell is from a parent processor, right three cells are possible configurations of children cells. *Top row:* 3D view, *bottom row*: corresponding 2D $\hat{y} - \hat{z}$ view.

3.5.3. Grid I/O

Structured meshes

Grid data from Chimera is output by default using the standard set of formats provided by the VTK library [90], and can be visualised and manipulated using ParaView, a software developed by Kitware Inc. [91]. The two VTK file formats used for logically rectilinear grids are listed here along with the additional data required to form the grid structure, i.e. excluding the variable data itself. Note, for a total of N_x cells in the $\hat{x}$-direction the number of cell vertices/points is $N_x + 1$.

- `.vti` format (VTK image data) - used for Cartesian grids with equally sized cells. This format requires the coordinate of the origin $[x_0, y_0, z_0]$, the spacing in each direction $[\Delta x, \Delta y, \Delta z]$, and the number of cells in each direction $[N_x, N_y, N_z]$, giving a total of six reals and three integers.

- `.vts` format (VTK structured data) - used for rectilinear but unequally-spaced, cylindrical or spherical, grids.
  This format requires the three arrays of point coordinates in each spatial direction x_p, y_p, z_p, giving a total of $N_p = (N_x + 1) + (N_y + 1) + (N_z + 1)$ reals.

For the same number of cells the second format will (almost always) produce a larger size file. For 1D simulation output the number of point coordinates stored in the `.vts` format will be larger than total number of cells, $N_p/N > 1$, however this fraction decreases with simulation dimensionality. Similar considerations are required when moving to unstructured meshes.

Unstructured meshes

Unstructured meshes require a different set of file formats as they do not have logically rectilinear connectivity between cells. The most appropriate format depends primarily on the refinement strategy used, but also on the specific set of user requirements. Consider the example of a 3D Cartesian simulation of the magnetic reconnection experiments described in §6. Without going into details of the system, an example of a simulation mesh refined to four resolution levels is shown in Fig 3.17. The total number of processors specified from the splitting algorithm is 2962, each with $50 \times 50 \times 50 = 125,000$ cells. Typically output variables include six scalar (density, electron temperature, etc.) and four vector (**B**-field, **E**-field , current density, velocity), giving 18 3D scalars. Three output formats for this mesh are then possible.

- A naive but straightforward approach would be to interpolate all grid quantities onto a full uniform grid at the highest resolution, allowing all data to be written in the `.vti` format. This would minimise the grid topology information to effectively zero, but also makes any memory savings from a variable mesh grid entirely redundant.

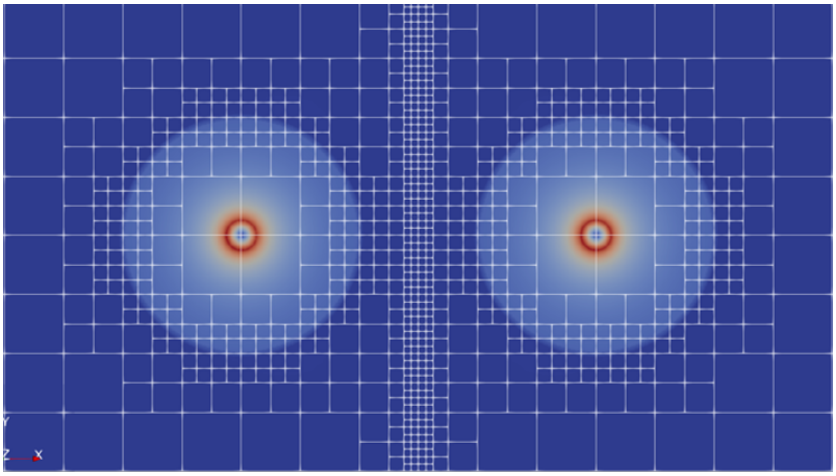


Figure 3.17.: Example of mesh splitting for a 2D simulation of the MARZ platform. A total of 2962 processors were used in this variable resolution configuration.

- In the block-based refinement approach, the advantage lies in using an unstructured grouping of regularly spaced blocks. Therefore, the `.vti` format can be used for data from each processor and stitched together to form the full grid. This is the approach adopted in *stare*, where output from separate `.vti` files are combined using a `.vtm` (VTK multi-block data) file. The grid information per `.vti` is considered negligible relative to the data, making this a reasonably inexpensive output. However, this file-per-processor approach for all ten variables results in $(2962 + 1) \times 10 \sim 3 \times 10^4$ files written per timestep. This places a high burden on the file systems of HPC facilities, and can result in slow visualisation software processing.

- A third approach can be adopted from cell-based refinement strategies, where the topology of the grid is arbitrary and therefore must be constructed from the point coordinates of each individual cell. For a grid of $N = 125,000 \times 2962 \approx 3.7 \times 10^8$ cubic cells, this requires an additional $8N \times 3 = 24N$ real values (3 spatial coordinates of 8 vertices) to fully define the grid. The `.vtu` format (for VTK unstructured data) supports this output structure, however additionally requires specifying an integer for each cell's type, with further connectivity information per cell. The result is a highly memory intensive output structure. However, for a static implementation the grid topology does not vary after initialisation, therefore a grid structure file could hypothetically be written once and referred to for all variable outputs at all timesteps. This is still more memory-intensive than the above method of using `.vtm` files, however would minimise the number of files required. The option of a separate grid topology file does not currently exist in the VTK library, and other libraries are being investigated as a future output option.

3.6. MHD variable transfer

The section extends the discussion in §3.5.2 to MHD variables. The difference between MHD and hydrodynamic variables are twofold; firstly, their difference in locations on the grid, and secondly the conservation law obeyed. Rather than be conserved in a volume-averaged manner the magnetic-field must instead remain divergence free, which is naturally enforced by use of the vector potential solver on a staggered Yee grid. However, this condition becomes more challenging to adhere to across differences in resolution. The methods in literature for treating these are described here first, followed by a description of the novel algorithm implemented in Chimera.

3.6.1. Review of literature

In his 2001 paper [11], Balsara describes schemes for prolongation and restriction of magnetic field, for a system that obeys:

$$\frac{1}{c}\frac{\partial \mathbf{B}}{\partial t} = -\nabla \times \mathbf{E} \tag{3.4}$$

$$\mathbf{E} = -\frac{1}{c}\mathbf{v} \times \mathbf{B}. \tag{3.5}$$

This is referred to as constrained-transport, which updates the magnetic and electric field variables only [92]. The divergence-free property of the **B**-field is preserved through use of a staggered Yee grid, in a similar fashion to the vector potential method (§2.1.3). The restriction operator is straightforward, and akin to the condition in Eq 3.1. A simple area-weighted (instead of a volume-weighted) average can be taken of the field components where the coarse cell faces coincide with finer cell faces (Fig 3.18, top). The exception is when components of the two grids are coincident, in which case the sum of the electric field components on the coincident edges should be taken (Fig 3.18, bottom). This still yields the area-weighted sum of the face-centred magnetic field, however introduces an extra correction to the other components, akin to the flux correction in Eq 3.2.

The prolongation procedure is somewhat more complicated. Here, the three magnetic field components are assumed to adopt a high-order polynomial profile within the volume, such that any point within the cell will be divergence free:

$$B_x(x,y,z) = a_0 + a_x x + a_y y + a_z z + a_{xx}x^2 + a_{xy}xy + a_{xz}xz + a_{xyz}xyz + a_{xxy}x^2y + a_{xxz}x^2z$$
$$B_y(x,y,z) = b_0 + b_x x + b_y y + b_z z + b_{xy}xy + b_{yy}y^2 + b_{xz}xz + b_{xyz}xyz + b_{xyy}y^2x + b_{yyz}y^2z$$
$$B_z(x,y,z) = c_0 + c_x x + c_y y + c_z z + c_{xz}xz + c_{yz}yz + c_{zz}z^2 + c_{xyz}xyz + c_{xzz}z^2x + c_{yzz}z^2y$$

By imposing the divergence-free condition for the above set of continuous equations, as well as the discrete version across the cell volume:

$$\frac{B_x^+ - B_x^-}{\Delta x} + \frac{B_y^+ - B_y^-}{\Delta y} + \frac{B_z^+ - B_z^-}{\Delta z} = 0, \tag{3.6}$$

where $B_x^{\pm} = B(x \pm \Delta x/2, y, z)$ and the same for y, z, the full set of coefficients can be obtained in terms of the cell and its neighbouring values. This has been implemented in various AMR-MHD codes with success [87] [86]. However, as acknowledged by Balsara himself, this formulation is not suitable for resistive solvers, which replace Eq. 3.5 with

$$\mathbf{E} = -\frac{1}{c}\mathbf{v} \times \mathbf{B} + \frac{c}{4\pi\sigma}\nabla \times \mathbf{B}. \tag{3.7}$$

This requires restriction and prolongation operators that are additionally *curl-preserving*. A later publication by the Balsara [93] introduces various curl-preserving schemes at increasing spatial order using a similar procedure as described above. However, the complexity of these expressions makes them challenging to rapidly test and verify. Additionally, the spatial degrees of freedom for each component is a requirement for adaptive codes, where children cells are produced from within parents cells and cannot introduce divergence or curl into the existing solution. As such, a bespoke solution for a static refinement capability in a resistive-MHD code was developed.

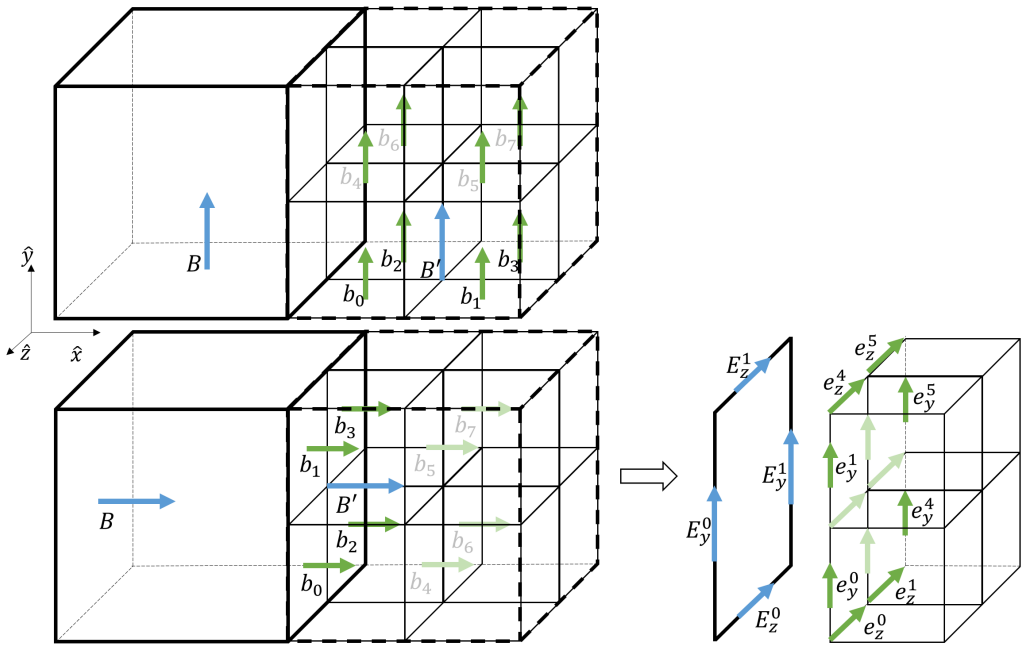


Figure 3.18.: Restriction procedure for magnetic field in Ref [11], for an interface between different resolution grids in the $\hat{x}$ direction. *Top:* Restriction of the B_y component, where $B'_y = 1/4 \sum_{i=0,3} b_i$ for equal area faces. Similarly applies for B_z component. *Bottom:* Restriction of the B_x component, where the low and high resolution grids share a face. The same area-weighting sum is applicable, however this is automatically achieved by replacing the coarse electric field values with an average of the finer electric field that share an edge.

3.6.2. Implementation in Chimera

Constrained-transport solver

As a precursor to introducing mesh refinement into the MHD routine, the existing vector potential formulation was first recast into a constrained-transport (CT) formulation. These are described in Algorithms 1 and 2 respectively, and for clarity the steps at which boundary conditions are applied are also mentioned. On inspection the algorithms are identical, however the latter was found to be more straightforward as the number of variables to keep track of in a self-consistent manner is reduced. While the CT solver no longer uses $\mathbf{A}$, an extra term called $\mathbf{E}_{\text{temp}}$ is introduced. As is evident from Algorithm 2, this term is distinct from the standard EM variable as it is not evolved over time; instead it is a temporary flux variable used only to update $\mathbf{B}$. This term could simply be injected into the expressions for $\mathbf{B}^n$ and $\mathbf{B}^{n+1}$ in Algorithm 2, thereby avoiding the need to store an additional 3-dimensional variable which further adds to memory constraints, however it was deliberately kept distinct in order to apply the following exchange procedure.

E/B exchange procedure

The exchange algorithm is based on the following two points:

- Solvers that use a sum of fluxes to advance a cell variable in time typically calculate both the left and right fluxes in each direction in every cell. This leads to a redundancy in calculation, as the right-hand flux of a cell is equal to the left-hand flux its neighbour on the right.
- The staggered Yee grid formulation requires that $\nabla \times \mathbf{B}$ is calculated using neighbouring cells that are to the left/bottom/back (referred to henceforth as lower faces) of the cell, while $\nabla \times \mathbf{E}$ is calculated using neighbouring cells that are to the right/top/front (referred to as upper faces) of the cell, see also Fig 3.19:
  - On the left, $\mathbf{E}_z^{i,j,k} \sim \nabla_z \times \mathbf{B}$ requires cells $[i,j,k]$; $[i-1,j,k]$, and $[i,j-1,k]$
  - On the right, $\mathbf{B}_z^{i,j,k} \sim \nabla_z \times \mathbf{E}$ requires cells $[i,j,k]$, $[i+1,j,k]$, and $[i,j+1,k]$,

  and extends to the Cartesian rotations $x/i \to y/j$, $y/j \to z/k$, and $z/k \to x/i$.

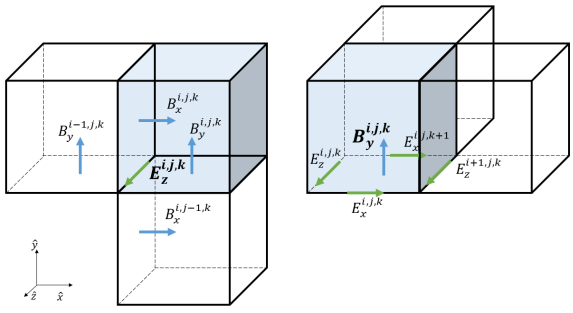


Figure 3.19.: Demonstration of asymmetric stencils for $\mathbf{B}$ and $\mathbf{E}$ on a Yee grid.

The exchange procedure is then described as follows. Note, **E** can either denote the electric field, or the effective flux term $\mathbf{E}_{\text{temp}}$. These four steps are carried out in Algorithm 2 every time there is a pair of $\mathbf{E}/\mathbf{E}_{\text{temp}}$ and **B** updates.

STEPS

1. Receive ghost cell values of **B** from lower faces (equivalently, send to upper faces)
2. Calculate **E** using $\nabla \times \mathbf{B}$
3. Receive ghost cell values of **E** from upper faces (equivalently, send to lower faces)
4. Calculate **B** using $\nabla \times \mathbf{E}$

These are demonstrated in Fig 3.20 in the context of sending **B** and **E** variables across processors with respect to P_1. For simplicity this is shown in 2D, with only B_x, B_y, and E_z components considered. P_1 receives magnetic field ghost cell values from its lower neighbours (white arrows), calculates **E** using the curl of **B** on its own grid, then sends back its lower neighbours (while also receiving electric field ghost cell values, shown in white circles, from its upper neighbours), and updates magnetic field using the curl of **E**. Therefore rather than perform a standard exchange across the overlapping region (described in §3.4.1), different sets of transfers are performed in the lower and upper directions. While this asymmetry may appear concerning, it is in fact just a way to mitigate redundant flux calculations over cell faces (first of the two points above).

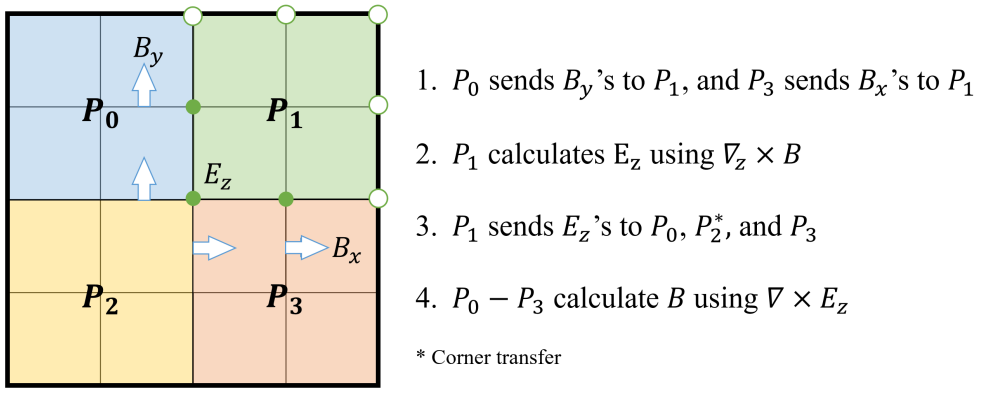


Figure 3.20.: Demonstration of **B**/**E** exchange routine for the updated magnetic field solver. P_1's ghost cells are shown in white arrows for magnetic field components, and white circles for electric field. Corner transfers are discussed in Fig 3.23. For clarity, only relevant domain cells that are involved in the transfer are shown.

Variable mesh boundary conditions

Contrary to most solvers that incorporate mesh refinement, the above procedure is carried out irrespective of the neighbour's resolution. Consider the two examples in Fig 3.21, which show prolongation and restriction operations for P_1's magnetic field ghost cells. Rather than prolongating all three components of the magnetic field using high-order polynomial fits, only a linear interpolation of the relevant components across a face is required. The ghost cell values in Fig 3.21(a) are then given by:

$$B'_x = \frac{4b_x - B_x}{3} \qquad B'_y = \frac{4b_y - B_y}{3} \tag{3.8}$$

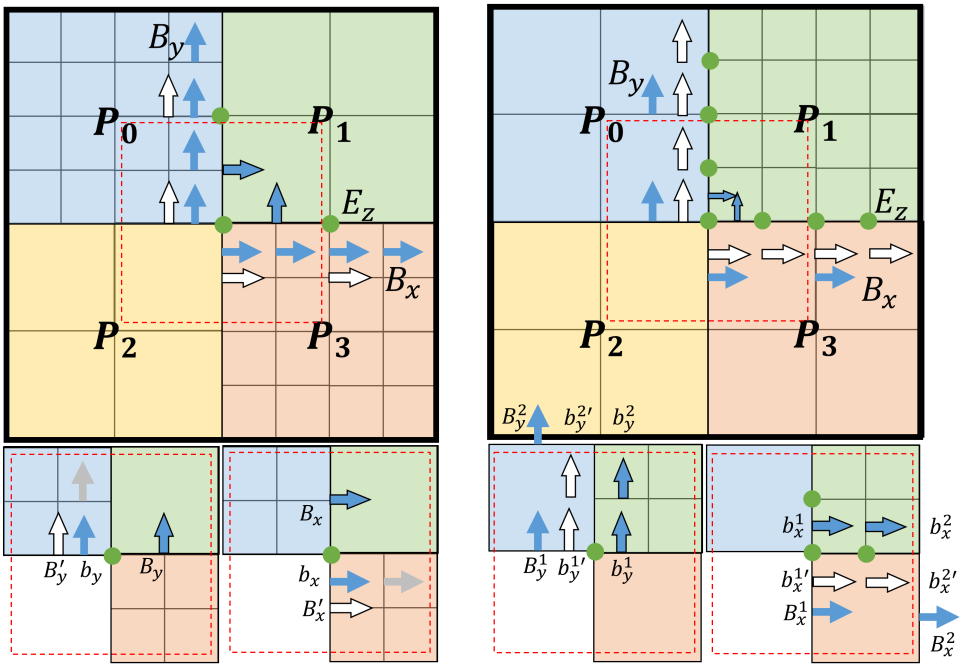
and in Fig 3.21 (b) are:

$$b_x^{1'} = \frac{b_x^1 - 2B_x^1}{3} \qquad b_y^{1'} = \frac{b_y^1 - 2B_y^1}{3} \tag{3.9}$$

$$b_x^{2'} = \frac{b_x^2 - 2(B_x^1 + B_x^2)/2}{3} \qquad b_y^{2'} = \frac{b_y^2 - 2(B_y^1 + B_y^2)/2}{3}, \tag{3.10}$$

where apostrophes indicate ghost cell values and lowercase/uppercase indicates high/low resolution cells. Note the additional averaging required for high resolution cells that lie in between low resolutions cells ($b_x^{2'}$ and $b_y^{2'}$). This method for restriction is different to the Balsara method described above, which took the area-weighted average of high resolution cells with cell faces coincident with low resolution cells, in order to preserve magnetic flux through the shared face $\Phi = \int_A \mathbf{B} \cdot \hat{n} dA = \sum_i b^i a^i$. Here the same coincident-face high resolution cells are used, however they are not averaged across the cell face. Minimal difference was observed in integrated test problems when using both methods, however further testing is required to quantify the difference.

A similarity with the Balsara method is the electric field restriction procedure, shown instead in Fig 3.22. Recall in this step the ghost cells belong to P_0 (blue processor), as **E** field components are sent from right to left. The restriction step on the right image shows that high resolution values are only used if their locations are coincident with the low resolution ghost cell edges (i.e. not the grey cell). In the case of prolongation on the left, the high resolution ghost cells either use a straight injection of the low resolution values $e_z^{1'} = E_z^1$, or take an average if they lie in between two, $e_z^{2'} = (E_z^1 + E_z^2)/2$.



(a) Restriction (b) Prolongation

Figure 3.21.: Variable resolution version of Fig 3.20, showing restriction and prolongation operations for **B** field components. Zoomed in subset of cells shown below with labelled values.

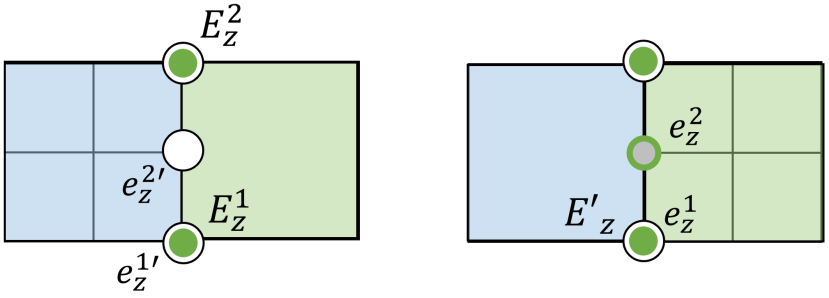


Figure 3.22.: *Left:* prolongation and *right:* restriction of **E** components for populating ghost cells (white circles) in P_0 (blue processor).

Extents

Before describing the extension to 3-dimensions, a few points about the extents are first discussed. In this implementation, the solver updates all three components of both the **B** and **E** variables in the range $i = [1 : n_x], j = [1 : n_y], k = [1 : n_z]$. The only two exceptions are 1) B_x is solved from $i = [1 : n_x + 1]$, and 2) the extents of E_y and E_z are increased to $i = 1 : n_x + 1$ if the processor lies on the upper $\hat{x}$ boundary (and the equivalent for other components). This is unlike the original solver, where electric field components are solved up to these extents by default.

These reduced extents mean every **B** and **E** component is uniquely defined on the grid (except for the B_x component at $i = n_x + 1$), as demonstrated for E_z in Fig 3.23. While this has its benefits, this means the corner cells are unaccounted for in the MPI communications (P_2 is missing information from P_1).

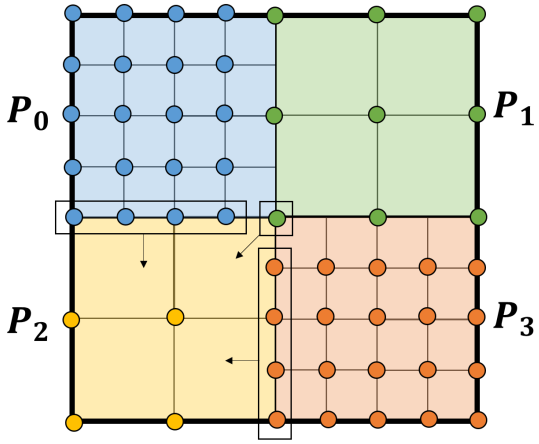


Figure 3.23.: Demonstration of how components of the MHD variables are uniquely defined at each grid cell, shown here in terms of E_z. Colour of dots corresponds to which processor the value 'belongs'. Black arrows indicate direction of communication.

Within this framework, this upper-right corner necessitates a separate communication. As a reminder, the original solver communicated corner transfers by extending the region being sent (Fig 3.5). This is straightforward in an unrefined mesh, as the same corner value will be communicated irrespective of whether an $\hat{x}$ or $\hat{y}$ sweep is being performed first. In the case of Fig 3.23, however, the value of E_z calculated by each of the processors P_0, P_1, P_3 at the centre-most edge will be different as each would store different values of B_x and B_y in their respective stencils. Enforcing a unique value at each grid point avoids potential inconsistencies, which can cause significant numerical instability. The computational cost of this extra transfer is discussed in further detail below, however overall it is relatively small to ensure a smooth solution without requiring complex prolongation.

Extension to 3D and variable splitting

The restriction and prolongation operators for the magnetic and electric fields can be trivially extended to 3-dimensions. For restriction operators, all high resolution values with faces and edges that are coincident with the corresponding low resolution ghost cell faces/edges are averaged. For prolongation, the high resolution coincident ghost cells use simple injection of the low resolution values (Eq 3.9), while the non-coincident cells use low resolution averages (Eq 3.10). The equivalent 3D diagrams of Figs 3.21 and 3.22 are shown in Fig 3.24, for B_y and E_z transfers in the $\hat{x}$-direction. Using standard Cartesian rotations, the full set of B_m and E_m in the n-the direction can therefore be deduced for $m, n = i, j, k$ and $m \neq n$.

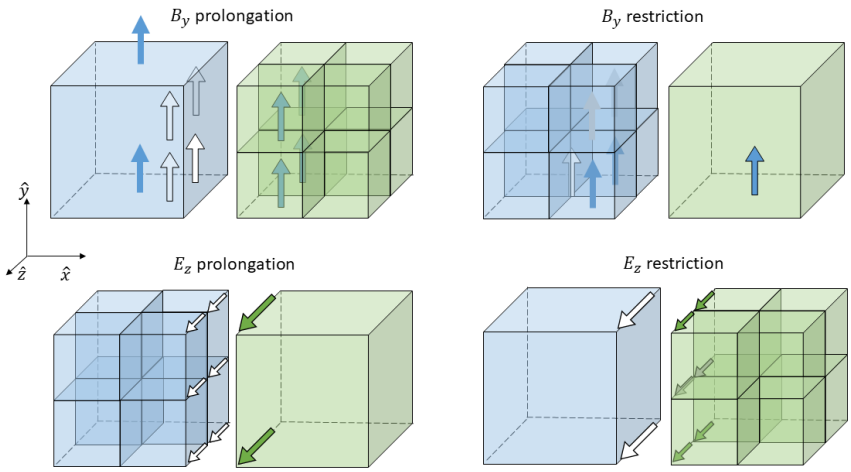


Figure 3.24.: 3D representation of prolongation and restriction operations for B_y and E_z. Simultaneous B_z transfers for $E_y \sim \partial_x B_z$, and E_y transfers for $\partial_t B_z \sim \partial_x E_y$, are not shown for clarity. Recall, ghost cells belong to P_1 (green processor) for B_y transfer, and to P_0 (blue processor) for E_z.

An additional layer of complexity in the implementation is introduced to the above method due the flexible choice of splitting described (see Fig 3.15 and associated discussion). This is largely encapsulated in the number of MPI communications and indexing when packing data to transfer, which has already been discussed in the context of transferring cell-centred variables. A further consideration is required for linear interpolation in prolongation and restriction, for example populating the ghost cells B'_y, is no longer described by Eq 3.8 but instead by:

$$B'_y = \frac{2x_m \bar{b_y} + (1 - x_m) B_y}{x_m + 1}, \tag{3.11}$$

where x_m corresponds to the refinement ratio, i.e. $x_m = 2$ if the $\hat{x}$ direction is defined as being split, and $x_m = 1$ otherwise. In the former case Eq 3.8 is recovered, and in the latter case $B'_y = b_y$, as with straight injection. $\bar{b_y}$ is likewise either an averaged quantity in the $\hat{z}$ direction or not, depending on if $\hat{z}$ is set to be split. Similar equations can be derived for all components in all directions, however this is both trivial and tedious,

and therefore not provided here.

Discussion of method

The method described above has been implemented and shown to work successfully on a series of test problems, as shown in the remaining chapters of this thesis. There are certain benefits to this method that others, including Balsara's, do not necessarily have and they are reiterated here.

- First is its conceptual simplicity; it is predicated entirely on the rule that magnetic field values are sent to the right, and electric field values are sent to the left, irrespective of neighbour resolution. Rephrased, magnetic field (variable) ghost cells are only populated on lower faces of a processor, and electric field (flux) ghost cells on the upper faces.

- This leads to the second advantage; because almost all $\mathbf{B}/\mathbf{E}$ values are only solved between $i/j/k = 1 : n_x/n_y/n_z$ (i.e. the extents of domain cells in a processor), the two variables are uniquely defined in each grid cell, as shown in Fig 3.23. The only exception is B_x in the $\hat{x}$-direction, which is defined at $i = n_x + 1$ on the left neighbour and $i = 1$ for the right neighbour, however using the above transfer of $\mathbf{E}$'s these coincident values are consistent.

- Thirdly, this method does not introduce divergence. The constrained-transport solver updates the magnetic field by taking a curl of the electric field, while the original solver equivalently takes the curl of the vector potential, the divergence of which is zero to machine precision. Simply transferring values of $\mathbf{E}$ is therefore similarly a divergence-free operation.

There are two main shortcomings of this method. Firstly, while enforcing the directional transfer of magnetic and electric field mitigates significant complexity, this enhances the inhomogeneous variations in spatial accuracy that already exist as a consequence of using variable resolution. Consider Fig 3.25, which shows a block of high resolution processors surrounded by low resolution processors. If the stencil for calculating electric fields $\mathbf{E} \sim \nabla \times \mathbf{B}$ (shown in dashed horizontal and vertical lines) uses low resolution cells then this will decrease the spatial accuracy. Performing the curl operation on high resolution cells with smaller values of Δx and Δy will increase the spatial accuracy. The result is a combination of higher and lower accuracy fluxes. Other mesh refinement codes typically enforce the prevailing flux value to originate from the higher resolution processors, however this not viable using the aforementioned exchange. This variation could be addressed by using higher-order reconstruction schemes and will be considered as a future avenue of improvement.

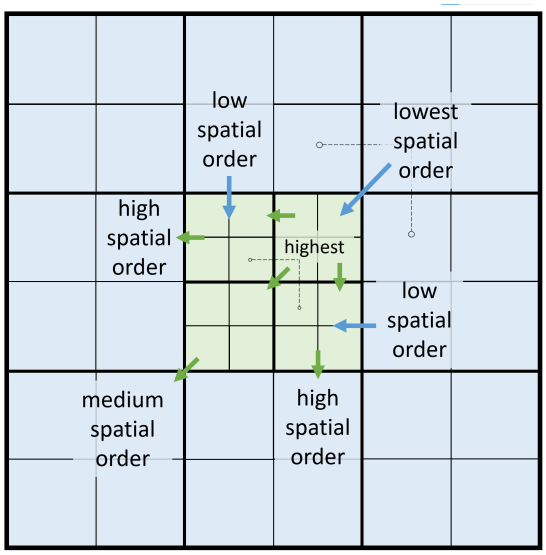


Figure 3.25.: Variations in spatial order of reconstruction schemes due to using a combination of high and low resolution cells in a stencil.

The second is the additional set of MPI communications due to corner transfers, as was described with respect to Fig 3.23 above. This is typically undesirable due to the latency from communication across processors, and is avoided in the standard communication using extended transfers in the three cardinal directions. Consider the transfer of variables in the original solver in the $\hat{x}$-direction only. This involves separately sending and receiving in the left and right directions the two variables B_y and B_z, i.e. eight transfers. In this method, the send/receive operation on a given face occurs simultaneously and is therefore considered one transfer (see §3.5.2). With that in mind, the total number of transfers are a left-receive and right-send of B_y and B_z, a right-receive and left-send of E_y and E_z, and an upper-right-receieve/lower-left-send (corner exchange) of E_z, i.e. five transfers. The communications in this new method are further optimised to combine the transfer of two components into one MPI call, reducing the number even further. In total, the original method uses 24 transfers in 3D dimensions to exchange MHD variables, whereas the new method uses nine. Therefore, the addition of corner transfers is not considered detrimental to performance.

4. Single wire Z-pinches

Several terms have been used to describe a solid-density wire driven by a high current discharge to form a Z-pinch, however in this work the term 'single wire' is used. Single wires have been fielded on several facilities in various configurations, including insulating wire materials, metallic wires, wires in vacuum, in insulating media, wires fielded on generators with long rise times, short rise times, etc. As the wire evolves it traverses a large region of thermodynamic phase space, making it a system of interest in many contexts. One such example is the use of cryogenic deuterium or deutrated carbon fibres to reach fusion-relevant conditions in hotspots formed along the fibre lengths, as was discussed in §1.

In this chapter, however, the focus is shifted to metallic wires. The reason for this is twofold:

- Firstly, higher-Z elements than deuterium would be expected to undergo stronger radiative cooling and therefore reach even further compressed states in the hotspot. Strong x-ray emission has been observed from 'bright spots' in mid to high-Z wires, such as Al and W wires, however the progression followed to reach these conditions is not fully understood.

- Secondly, the 2-wire X-pinch and wire arrays (described in §5.2 and §6 respectively) are constructed from multiple mid-Z metallic wires. The current per wire in both is close to that experienced by a single wire on a smaller generator, therefore the dynamics of the latter will inform understanding of the former, excluding the effects of a global magnetic field.

The ability to simulate the hotspots formed in exploding single wires at sufficient resolution to capture the full range of length scales has previously been limited by computational resource, and therefore only modelled in one and two dimensions. Separately, the coronal plasma formed around metallic wires expands rapidly to large radii, requiring a large computational domain and therefore restricting the resolution at which the system can be modelled using a uniform grid. The work in this chapter presents the first systematic study of the single wire system from 1D to 3D at 1 μm maximum resolution using the mesh refinement capability described in the previous chapter. The dynamics in the one dimensional case are simpler, and therefore can be used to study the effects of varying physics models. Introducing higher dimensions adds increasing complexity, particularly the instabilities present (2D introduces the effect of $m = 0$ instabilities, 3D introduces $m = 1$ instabilities).

In order to provide context for later results, the chapter begins by reviewing literature of single wire explosions in vacuum in §4.1. The dynamics are described in terms of the formation of the core-corona structure, the ablation mechanism that sustains it, and each of the two regions in further detail. With the expected behaviour in mind, an overview of the specific experiment on which the simulations are based is provided in §4.2. The relevant diagnostics are described, as well as key findings used for later comparison against simulation. The simulation results are then presented in order of increasing dimension. The 1D results presented in §4.3 explore sensitivities to different model options, particularly equation of state, transport, and radiation models. The 2D results presented in §4.4 demonstrate the effect of including axial motion, which has consequences both for modulations in the coronal plasma and for the compressed state at pinch points. Finally, the 3D results presented in §4.5 assess the effect of azimuthal variations on the collapse process.

4.1. Review of dynamics

This section provides an overview of single wire behaviour, focusing specifically on metallic wires that explode in vacuum, relevant to the aluminium wire experiments in the following section.

Formation

Given the simplicity of the system one might expect to be able to use simple energy balance arguments to predict its evolution. For example, consider a 15μm Al wire experiencing a current pulse with 65ns rise time. The skin depth at this frequency is:

$$\delta = \sqrt{\frac{2\eta}{\mu_0 \omega}} = 40\mu\text{m} >> r_{\text{wire}}, \tag{4.1}$$

where η is the material electrical resistivity and ω is the EM wave angular frequency, approximately given by $\omega \sim \pi/2\tau_{rise}$. Neglecting effects such as the temperature dependence of η, Eq 4.1 implies that the field diffuses through the wire on timescales smaller than hydrodynamic processes. The wire might therefore be expected to undergo bulk Ohmic heating, phase transitions through melt and vapourisation phases, and finally uniform volumetric expansion.

Several radiography and x-ray backlighting experiments [16] [94] [95] [96] [17] have demonstrated this not to be the case. Instead what forms is a 'core-corona' structure, consisting of a cold, dense core in a multiphase liquid-vapour state, surrounded by low-density coronal plasma that carries majority of the current, shown in Fig 4.1. The mechanism for initiating coronal plasma formation is challenging to study and remains debated, however the leading theory is as follows. The wire is initially heated either by a pre-pulse or a low current in the main pulse, heating the wire and causing low-Z impurities trapped in the metal to be evaporated or desorbed [16]. These are ionised

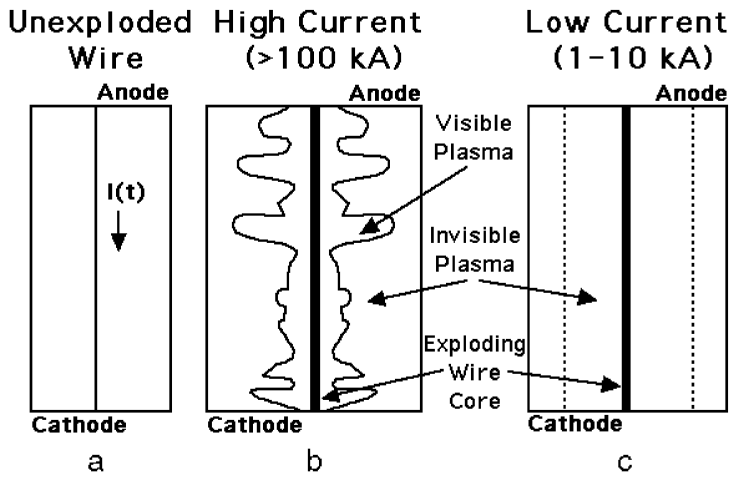


Figure 4.1.: Cartoon of wire ablation dynamics for high current and low current wires. In both cases, a low density coronal plsama forms around a high density core. The low current wires undergo a uniform explosion, while high current wires undergo visible $m = 0$ instabilities. Reproduced from Ref [12].

via thermionic emission or breakdown due to the applied voltage, initiating the coronal plasma formation.

Evidence for this process has been observed by varying the radial electric field polarity between positive and negative, in single wire explosions [97] and wire array experiments [98]. The negative polarity configuration in both cases were seen to undergo suppressed wire expansion rates. This was attributed to electric field expelling electrons away from the wire, reducing the rate of energy deposition and therefore expansion. Kinetic effects such as electron emission are not included in MHD models, therefore this process cannot be included in simulations. Instead, a low-density, high-temperature layer of gas is typically initialised around the wire core to recreate the effect of a low-Z gas discharge. While this has been used with success in the past, simulation results in §4.3 and §4.4 show how use of appropriate equation of state and transport models can reach equivalent results.

Ablation

Once the core-corona structure is formed, it typically persists for a significant fraction of the current pulse. This longevity is due to the corona carrying most of the current and therefore experiencing stronger Ohmic heating, which in turn heats and steadily ablates the core surface until a significant fraction of its mass has been depleted. The dynamics of energy transfer from corona to core, and mass transfer in the opposite direction, has been modelled by several authors; one of these models is described here.

A numerical study by Yu et al [13] uses 2D $r - \theta$ resistive-MHD simulations to study the dominant energy transfer mechanisms in the steady-state ablation of wire arrays. The mechanism for accelerating coronal plasma from stationary at the wire cores to super-Alfvénic at the array centre was investigated in terms of three distinct

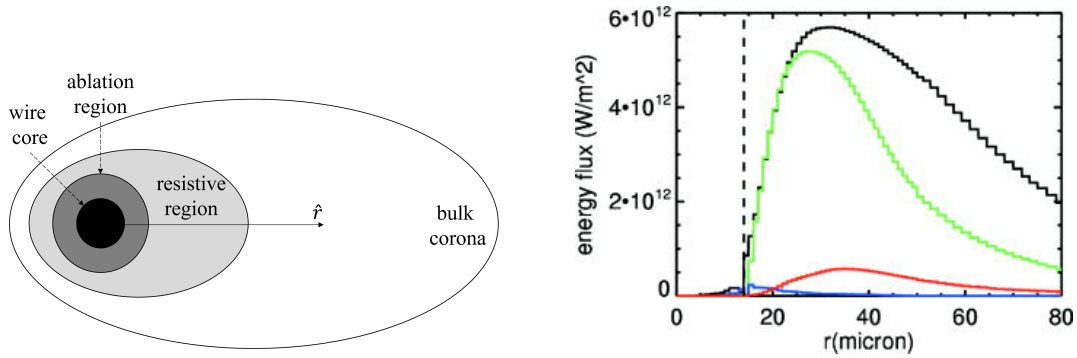


Figure 4.2.: *Left:* Illustration of distinct coronal plasma regions formed around the wire core in a wire-array (not to scale); ablation region (dark grey), resistive region (light grey), and bulk coronal plasma accelerated towards array axis (white). *Right:* Magnitude of the integral of fluxes plotted as a function of radius for different terms of the energy equation; radiative (black), convective (green), kinetic energy (red), and thermal conduction (blue). Dashed line indicates wire core radius. Reproduced from Ref [13].

coronal regions; 1) the bulk corona, 2) the resistive region, and 3) the ablation region, illustrated in Fig 4.2(a). The dominant energy fluxes through each region were identified by plotting surface integrals of the convective, conductive, and radiative terms in the energy equation (Eq 2.10) over the individual regions. Fig 4.2(b) shows the magnitude of these terms plotted as a function of radius for the ablation region, which extends up to ∼30 μm past the core. This shows the magnitude of the radiative heating term (black line) being comparable to the magnitude of the convective, cooling term (green line), while the thermal conduction contribution (blue line) is significantly smaller. While a similar analysis is beyond the scope of this work, the importance of the radiative term on the core/corona dynamics in 1D simulations of single wires will be discussed in §4.3.

Coronal dynamics

The above description of wire surface breakdown and continued mass ablation from core to corona pertain to both single and multi-wire systems. However, there is significant deviation between the two at later times due to the presence of a global magnetic field in the latter that redirects coronal plasma away from the standard radial compression it would otherwise experience (see discussion in §6.2). Additionally, the difference in field topology affects how the $m = 0$ instability forms. In both systems, the coronal plasma produced along the wire lengths forms modulations where $\mathbf{k} = k\hat{z}$ (where k is the modulation wavenumber), which feed into the imploded state at later times. In wire-array experiments the wavelength is found to stay roughly constant over time [99]. Radiography and laser-imaging were used to show that these "fundamental modes" scale as approximately $ka \sim 1$ (where a is the wire core radius), for example aluminium wires exhibit $\lambda = 0.5$ mm modulations with $ka \sim 1.6$, and tungsten wires $\lambda = 0.25$ mm with $ka \sim 1.3$. The modulations were demonstrated to be fundamentally MHD in nature rather than a consequence of wire material [100], and therefore have been described as

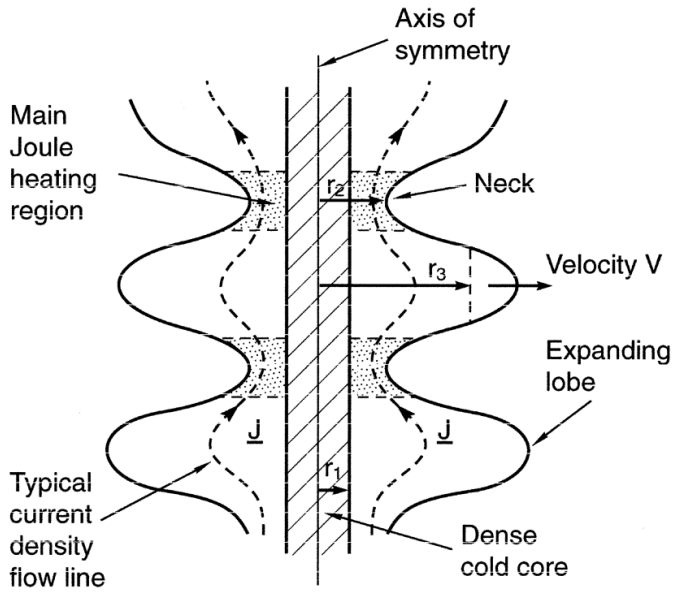


Figure 4.3.: Illustration of the $m = 0$ instability, showing central, uniform core and low density plasma forming neck and lobe regions. Preferential current path shown in dashed line. Reproduced from Ref [14].

a "modified $m = 0$" mode [101].

The dominant wavelength of a 'standard' (cylindrically symmetric with respect to the wire) $m = 0$ instability changes with time, and was described analytically in Ref [14]. In this model, the Z-pinch consists of a dense core surrounded by coronal plasma that has formed a series of necks and expanding lobes, as shown in Fig 4.3. The current density peaks in the neck regions where strong Ohmic heating occurs, while the lobes are sufficiently resistive that the current is not transported radially outwards with the expansion. In the early, linear-growth phase of evolution the radial displacement ξ_r from its initial radius ξ_0 is given by [14]:

$$\xi_r(t) = \xi_0 \exp\left[(ka)^{1/2}\frac{c_s t}{a}\right], \tag{4.2}$$

where c_s is the sound speed, suggesting shorter wavelength modes have the fastest growth at early times. However, these are also shown to saturate earlier. At later time when the behaviour becomes non-linear, ξ_r can be shown to peak at a wavelength of $ka = 1/\beta$ (Fig 3 in Ref [14]), where β is the ratio of thermal pressure to magnetic pressure. Excluding the lobe regions that are free to expand at the sound speed, the rest of the Z-pinch experiences a magnetic pressure comparable to the thermal pressure i.e. $\beta \sim 1$, which is consistent with experimental observations.

Core dynamics

Once the core is initially heated it continues its expansion at roughly the sound speed, whilst being steadily depleted from its outer surface through convective transport of mass into the corona. The sound speed at the core surface and rate of depletion are

slow relative to supersonic, Alfvénic processes in the corona, and can therefore be treated as relatively stationary over the pulse rise time. The state of the core during this period has been extensively studied using high resolution X-pinch radiography, which reveal fine scale structures at resolutions that optical imaging cannot access [16][12]. These experiments were performed using a relatively low current and long rise time pulser (1-4.5 kA in a 350 ns rise time) to mimic the prepulse on larger generators. These revealed a foamlike structure within the core which constitute an inhomogeneous liquid-vapour mixture, as shown in Fig 4.4. Horizontal striations were also often observed with scale lengths of the order of ~ 10 μm, which are particularly visible in Fig 4.4(a). At much later times (several μs) after peak current the core appears to coalesce into large droplets (> 10 μm) or threadlike structures as shown in Fig 4.4(b), suggestive of the core reaching a metastable state.

The origins of the often observed axial striations remain unclear, however these are thought to be seeded by the electrothermal instability (ETI). The ETI is also hypothesised to be the seed for magneto-Rayleigh Taylor (MRT) instabilities that occur on MagLIF liner surfaces and disrupt the stagnated fuel column, detrimentally impacting fusion yield [39]. This instability is a feedback loop whereby resistive inclusions in a material can cause uneven heating depending on its regime in conductivity space, leading to the formation of banded structures. In the metallic phase where $d\sigma/dT < 0$ the ETI results horizontal striations, and in the plasma phase where $d\sigma/dT > 0$ the ETI results in vertical filaments [102].

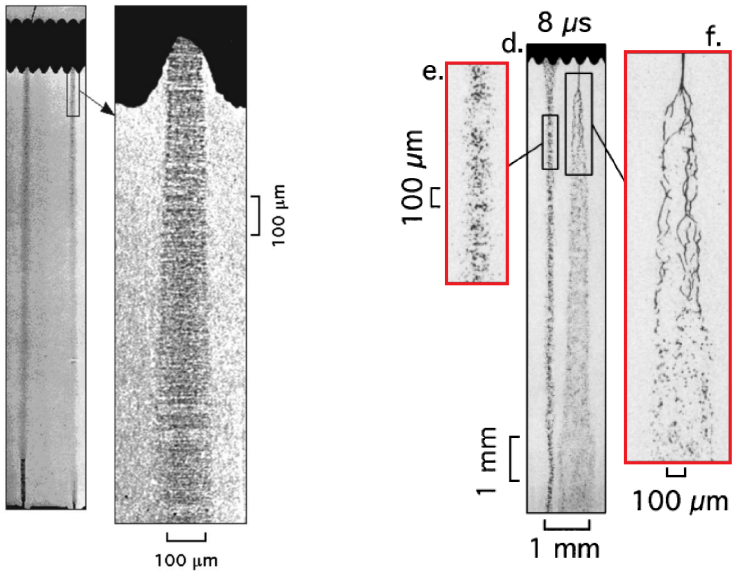


Figure 4.4.: *Left:* Radiograph showing horizontal striations formed across a 100 μm wire, reproduced from Ref [15], *right:* radiograph showing single wire metastable liquid-vapour state resulting filaments and bubbles. Reproduced from Ref [16].

4.2. Experimental overview

An investigation into early time single wire dynamics was conducted at Imperial College London in 1999 [17]. A set of varying diameter tungsten wires and $15\mu m$ aluminium wires were fielded on the IMP generator, with 150 kA peak current and a 65 ns rise time. This magnitude of current is comparable to that carried by individual wires in an array configuration on a larger generator (e.g. eight wires on the 1.4 MA MAGPIE facility at Imperial College [103]), making the generator a useful test bed for wire array studies. Additionally, 1D and 2D MHD simulations of the aluminium experiments were performed close to the time to support experimental findings [18]. The simulations are revisited over two decades later using latest code improvements, and it is these recent studies that the remainder of this chapter focuses on.

An extensive suite of diagnostics were fielded in the experiments. However, only results from Schlieren imaging and radial streak imaging are discussed here, as they provide core/coronal growth rates and modulation wavelength information used for comparison against simulations in later sections (§4.3 and §4.4). Each are briefly described, and relevant observations are highlighted.

- *Laser Schlieren imaging:* This technique uses a laser backlighter to detect gradients of refractive index along the line of sight. Low density regions (such as coronal plasma) result in small deflections and therefore appear as regions of bright emission. High density regions (such as the core) are strongly deflected and therefore detection is 'inaccessible', resulting in low emission. Eight images of separate experiments were taken at different times, between 8 ns and 97 ns in roughly 10 ns increments, shown in Fig 4.5. Plasma discharge is observed as early as 8 ns. A core-corona structure evident by 24 ns, although the spatial resolution is too poor to distinguish between separate regions before this time. Formation of $m = 0$ instabilities are visible from early time, and the wavelength of these modes are seen to steadily increase. By 56 ns, the instability necks have penetrated the core, breaking it up into 'islands' of high density. By 67 ns, the instability wavelength

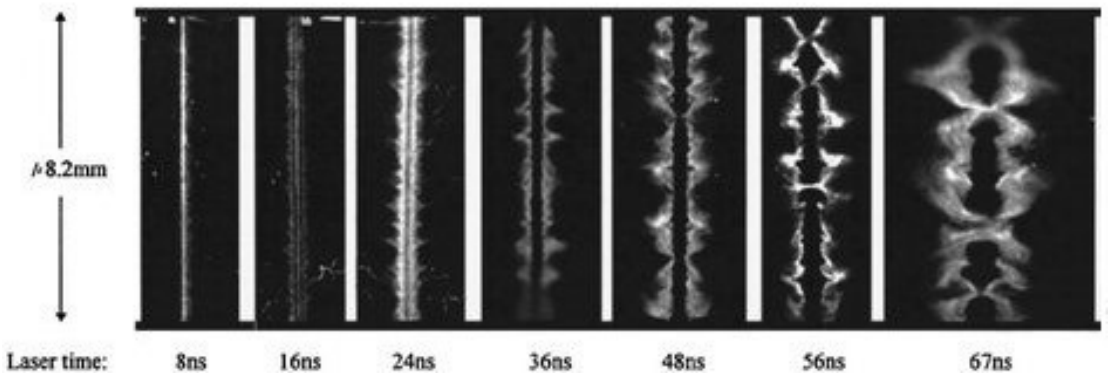


Figure 4.5.: Time series of Schlieren images of 15μm Al wire explosion in vacuum. Reproduced from Ref [17].

is of order $1-2$ mm, similar to the core width. The inaccessible, core regions were calculated using this time series to expand at a roughly constant velocity of 0.72 cm/μs.

- *Optical streak imaging:* Radial streak images of the plasma's optical self emission were obtained using a 300 μm wide slit to image one section along the wire's length at $1-2.25$ ns temporal resolution, as shown in Fig 4.6(a). The image shows a continuous expansion in the radial direction, from which the highest intensity emission was calculated to expand at $2.2 \pm 0.2 \times 10^4$ m/s. No re-compression is observable, which suggest the axial location used may have corresponded to the expanding bulge of an $m = 0$ instability, rather than a neck region which pinches at later times as seen in Fig 4.5.

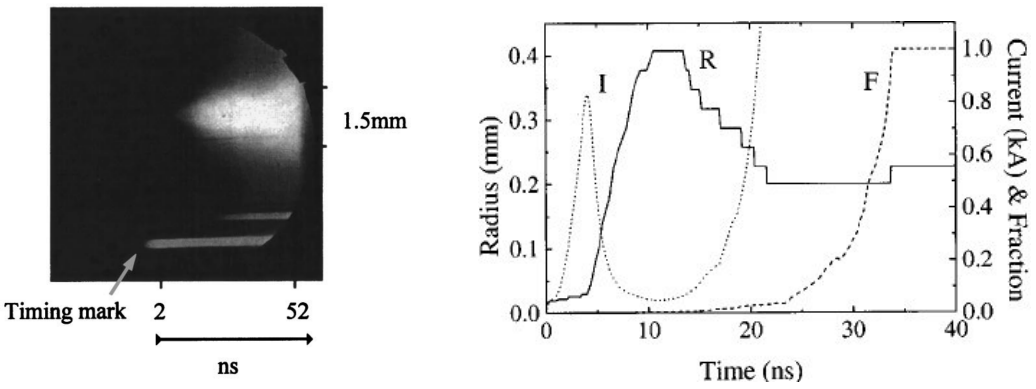


Figure 4.6.: *Left:* Optical streak imaging of 15 μm Al wire explosion. *Right:* 1D simulation of wire explosion on the left, showing radius of coronal plasma in solid line, current in the core, and fraction of the wire not in a condensed phase in dashed lines. Reproduced from Ref [17] and Ref [18].

Previous simulations of experiment

Before discussing recent simulation studies of the above system, results from an earlier numerical study are briefly described. Notable features of this version of the code include an implicit resistive diffusion solver, vacuum cut-off density of 10^{-4} kg/m^3, modified Thomas-Fermi equation of state to account for near-zero pressure at solid density, Lee-More transport coefficients, and a self-consistent circuit model to provide the current profile. The evolution of the 1D simulations is shown in Fig 4.6(b), and can be described as follows. A dwell phase exists for the first few nanoseconds, while the core is being Ohmically heated but remains in its condensed phase. Rapid surface expansion occurs until 10 ns, by which time most of the current has switched to the corona. Once the magnetic field is sufficiently strong, the low density plasma begins to re-pinch, i.e. contract towards the axis. The field is advected with the fully ionised plasma, and an increasing fraction of the current switches into the core, raising its temperature and therefore pressure to $\sim$ 3 GPa. This behaviour is compared against results from an improved version of the code in §4.3.

4.3. 1D simulations

This section revisits simulations of the aforementioned experiments using Chimera's upgraded capabilities, including improvements to EoS, transport, and radiation models. 1D axisymmetric simulations were run on a domain with $L_x = 1$ mm and a resolution of $\Delta r = 1$ μm. The same profile used to approximate IMP's current (measured by B-dot probes) as in Ref [18] is used here:

$$I(t) = I_{\text{peak}} \sin^2\left(\frac{2\pi t}{4 \times \tau_{\text{rise}}}\right), \tag{4.3}$$

where $I_{\text{peak}} = 150$ kA and $\tau_{\text{rise}} = 65$ ns. As in previous studies a cut-off vacuum density of $\rho_{\text{vac}} = 10^{-4}$ kg/m^3 is used, to prevent the simulation timestep from becoming too restrictive due to large Alfvén speeds in low density plasma.

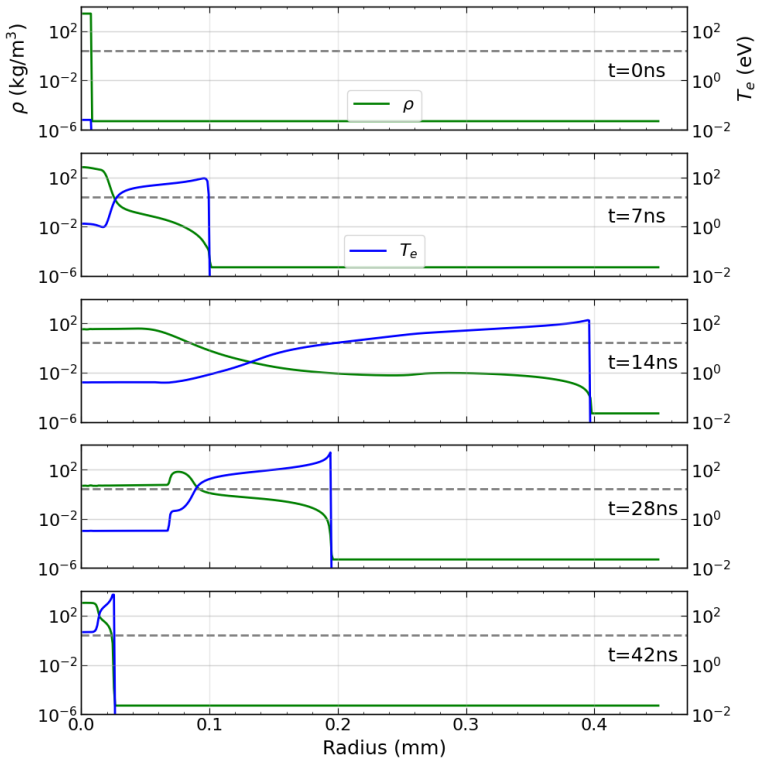


Figure 4.7.: Evolution of 1D single wire behaviour, showing density (green) and electron temperature (blue) profiles at 5 different times. The core-corona divide forms by 7 ns, and the corona reaches peak expansion at 14 ns, after which it reverses back onto the core. Regions of constant ρ correspond to vacuum.

The broad dynamics of this system is shown in Fig 4.7. The initial state of the 7.5 μm-radius wire is aluminium solid density at room temperature ($\rho_0 = 2.7 \times 10^3$ kg/m^3 and 293 K), as seen in the top-most plot. After only a few nanoseconds, the density profile shows a relatively dense region extending 20 μm, with a lower density region extending a further 800 μm. The temperature profile is inverted, with the 1 eV dense region surrounded by 100 eV plasma. These two distinct regions are identifiable as the core and corona respectively. The threshold lies at roughly 2.7 kg/m^3 = $\rho_0/1000$,

shown in the dashed grey line. By 14 ns, the corona has expanded rapidly, while the core undergoes a gentler expansion. This is near the point of peak radius, after which the corona begins to re-compress towards the axis via the Lorentz force. The compression results in PdV heating of the ions, which in turn heats the electrons as evidenced by the peaked edge in electron temperature at 28 ns. Between 20 ns either side of peak current, the entire mass is compressed and confined to the axis at $r \sim 30$ μm and $\rho \sim \rho_0$. Finally, the current decreases sufficiently to allow the wire to expand freely.

This behaviour is shown at higher temporal resolution in Fig 4.8. The radial streak profile in terms of pressure is shown with a time resolution of 0.1 ns along the y-axis, revealing more structure in the dynamics. For example, after 10 ns (during the compression phase) a high pressure shock is formed in the corona, which undergoes multiple reflections before the final strongest compression. Additionally, the pressure profile in the core between ~10-30 ns appears speckled as the core states move between finite and non-zero pressures due to falling in and out of the binodal curve (see Fig 2.6), which agrees with the formation of liquid-vapour states inferred from experimental radiographs in Fig 4.4.

By overlaying the current profile (Eq 4.3, shown in dashed black line) it is possible to see when the field strength and therefore magnetic pressure is strong enough to overcome the thermal pressure. In this simulation, the wire retains sufficient thermal pressure at peak compression to prevent the onset of radiative collapse as evidenced by its near constant radius. While these streak plots are highly informative, they are less useful for model comparisons as performed in subsequent sections. Instead, the interface profile

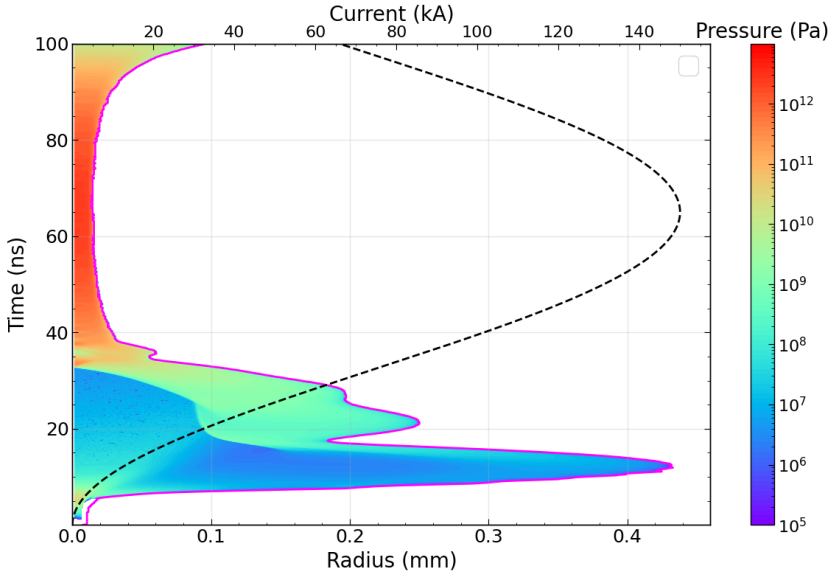


Figure 4.8.: Streak image of the pressure profiles at 0.1 ns temporal resolution and 1 μm spatial resolution. Overlaid around streak image is the interface profile, shown in magenta to represent a specific set of simulation model configurations. Dashed black line indicates current pulse used (upper y-axis).

of the plasma-vacuum boundary (magenta line here) is used. Each colour corresponds to a different configuration of model options, listed in Table 4.1 for reference.

Configuration	Colour	EOS	Transport	Radiation
I	Red	TF	Spitzer	Simple loss
II	Blue	FEOS	Spitzer	Simple loss
III	Magenta	FEOS	LMD	Simple loss
IV	Yellow	FEOS	LMD	No loss
V	Green	FEOS	LMD	Group-resolved loss
VI	Black	FEOS	LMD	Transport

Table 4.1.: Model configurations for 1D simulations. Colours correspond to interface profiles for Figs 4.8-4.11.

4.3.1. Effect of equation of state

This section studies the sensitivity of ablation dynamics with respect to the equation of state used. As mentioned above, past simulations from Ref [18] performed similar analysis using a modified Thomas-Fermi (TF) EoS to include condensed matter corrections. This model is compared here to the standard TF model without corrections, as well as the higher fidelity FEOS model.

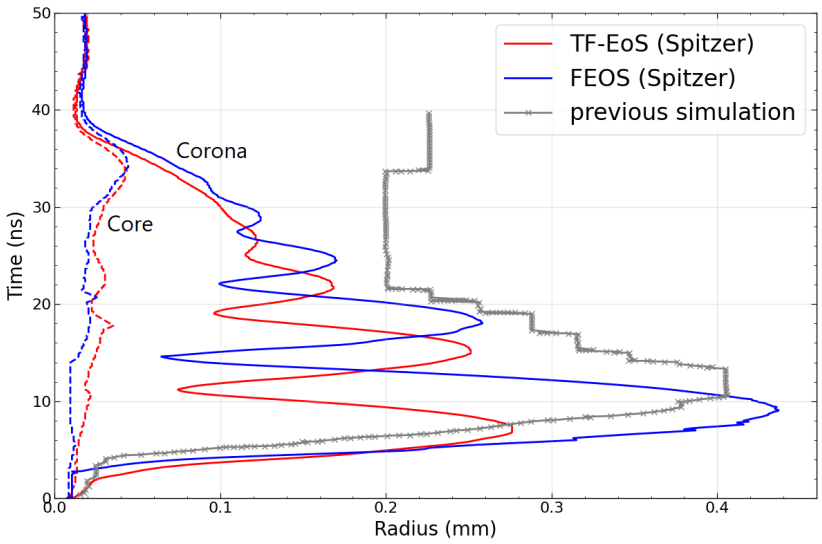


Figure 4.9.: Comparison of equation of state models using interface profiles for TF-EoS (red line), FEOS (blue line), TF-EoS with condensed matter corrections, run in previous simulation study (grey line) from Ref [18]. Solid lines correspond to the coronal plasma-vacuum interface, dashed lines correspond to the core-corona interface.

Fig 4.9 shows the difference in evolution across the three different equations of state, using the same Spitzer conductivity model and initial conditions. Note, solid lines indicate the coronal radius, which is the last plasma cell bordering the vacuum region, and dashed lines indicate the core boundary at $\rho = 2.7$ kg/m^3. There is an evident offset

in the timings, which originates from the beginning of the simulation. The standard TF-EoS (red line) predicts a pressure of $\sim$ GPa at ambient conditions for solid densities (see §2.3), and therefore begins its expansion at $t = 0$ ns. This drops the density and temperature of the core as it rarefies, altering the conditions of the corona. The Spitzer conductivity scaling $\sigma \propto T^{3/2}$ results in a reduced conductivity and therefore lower current through the corona, reducing the energy deposition and subsequent expansion.

Conversely, the FEOS simulation begins at 'zero' pressure at ambient conditions and is stationary until it is sufficiently heated to finite pressures (a few nanoseconds). This 'resistive phase' is also observed in experiments on similar timescales, and allows relatively more current to be shunted through the corona, forcing a stronger initial expansion as compared to TF-EoS. The next few cycles are relatively similar between the two lines, delayed only by the resistive phase offset. Finally, the compression in both cases converge to the same radius of 20 μm.

Comparing this against previous simulations (grey line) the presence of a dwell phase followed by a strong expansion bears closer resemblance to the FEOS results. However, the final compressed region reaches a much larger radius than in these results. While this could be due to differences in radiative models (see section below), it is also possible that the self-consistent circuit model used in previous simulations accounted for changes in inductance of the wire, and therefore did not drive the system to the same compressed state. The presence of oscillations in simulations from this study are a physical effect that arise due to changing balance between magnetic and thermal pressure in the corona. It is unclear why this behaviour is not observed in the previous study.

4.3.2. Effect of transport model

Further differences arise when swapping the Spitzer conductivity model for the tabulated LMD option, while using the same equation of state. For example, the resistive phase (initial Ohmic heating to finite thermal pressures) is extended even further, and offsets the initial dynamics by an additional 4-5 ns. The first compression after expansion also does not reach as strongly compressed a state as using Spitzer. There are fewer subsequent bounces, and the corona largely remains expanded before the late time compression. Most noticeably, however, the rates of core expansion differ strongly. The expansion appears agnostic to the equation of state used as seen in Fig 4.9, as the core remains under 40 μm for the first 35 ns for both simulations, giving an average rate of 0.15 cm/μs. When using the LMD conductivity, however, it occurs at roughly 0.8 cm/μs, reaching a maximum core radius of 110 μm. The 'inaccessible regions' in experimental schlieren images, which are assumed to be the high density core, were observed to expand at 0.72 cm/μs as described in §4.2. Assuming the inaccessible region corresponds to the 2.7 kg/m^3 threshold defined here, core expansion rates in simulations run with the LMD model clearly show better agreement to experiment as compared to those run with Spitzer transport.

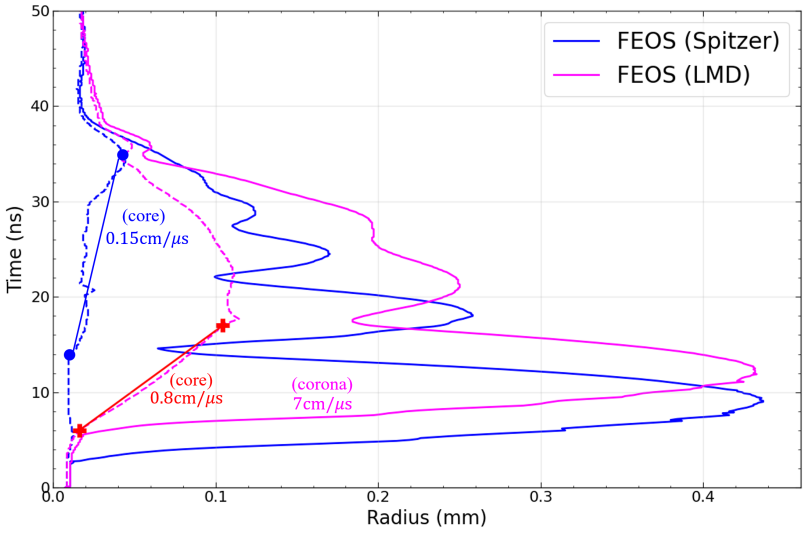


Figure 4.10.: Comparison of transport models using interface profiles for Spitzer (blue line) and LMD (magenta line). Gradient used to estimate rate of expansion of core regions for each model shown in dashed lines. Rates of core and coronal expansion labeled.

4.3.3. Effect of radiation models

This section assesses the effect of varying radiation models on the system dynamics. As demonstrated in the section above, the wire traverses a large region of thermodynamic phase space during its evolution, and therefore does the same across opacity space. While using radiation transport with high resolution atomic data tables is desirable to capture this variation, it is computationally prohibitive for large simulations. It is therefore worthwhile to understand the limitations of simpler models in reduced simu-

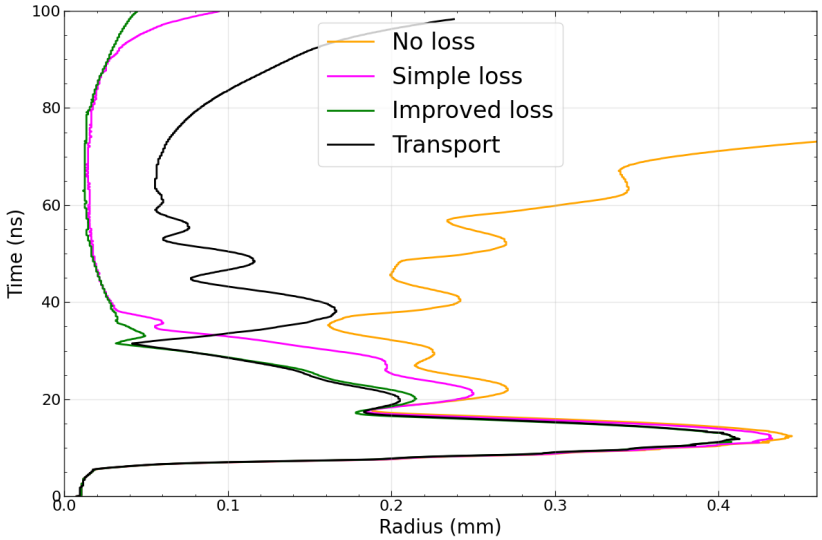


Figure 4.11.: Interface profile of plasma-vacuum boundary, comparing four different radiation treatments:no loss (yellow), simple loss (magenta), group-resolved or 'improved' loss (green), and radiation transport (black).

lations, and bear them in mind when utilising them in larger simulations. Note, atomic data tables with 32 spectral groups were used in both group-resolved loss and radiation transport cases.

Fig 4.11 shows the interface profile for the single wire experiments, comparing four different radiation models; no loss, simple loss, group-resolved loss, and radiation transport. As a reminder, previous simulations shown in red, blue, and magenta lines used the simple loss model (see Table 4.1). The dynamics of the magenta line (config. III) have been described in §4.3.1 and §4.3.2 therefore are not repeated. Only results from the no loss, group-resolved loss, and radiation transport are discussed. Note, the first two are compared against the latter, which is assumed to be the reference case.

No loss: It is evident from Fig 4.11 that entirely neglecting radiative losses does a poor job of capturing the reference case dynamics. The timing of the initial expansion up to 10 ns is well matched, however this is expected given the core is heated to only a few eV. By 10 ns, the corona has heated sufficiently (>50 eV) to begin strongly radiating. All models show a drop in coronal temperatures due to PdV expansion, however without an additional radiative cooling mechanism, the corona in the no-loss case continues its expansion for an extra 30 μm. Material closer to the axis that initially underwent Ohmic heating retains its temperature, such that by the time the corona begins to re-compress it encounters a higher thermal pressure and contracts only to 180 μm (as opposed to < 40 μm in the other cases). This larger coronal radius therefore encounters a lower magnetic field, and does not experience a sufficiently high compressional Lorentz force to counter its thermal pressure. A few, small amplitude, bounces occur due to an increasing current, but overall the system does not reach the minimum radius required to experience confinement, and expands rapidly out of the domain.

Improved loss: Including opacity effects in the loss model as described in §2.4 results in visibly better agreement to the early time dynamics of the reference case, particularly between $12 - 30$ ns. This indicates enhanced emission from line radiation in the corona at these conditions, which a simple continuum model excludes. The timing of the compression is well matched to the transport case, which the simple loss case overestimates by 8 ns. At 30 ns, however, there is significant departure between the group-resolved loss and transport case. The group-resolved case is confined against the axis between $13 \leq r \leq 40$ μm until 20 ns after peak current, showing closer match with the simple loss case, after which the compression is relaxed and the lines diverge.

Radiation transport: Closer inspection of the radiation transport case at $t \geq 25$ ns reveals the presence of a radiation front that travels from the high-temperature edge into the high density core. This moves radially inwards faster than the contracting plasma edge, and heats the inner plasma. Therefore the incoming plasma then rebounds on heated plasma, instead of being confined like the loss cases. What then follows is a series of further bounces for the transport case towards peak compression at peak current, before the final expansion. The loss cases reach a minimum radius of 13 μm, while the transport case reaches 55 μm. This considerable difference can be understood using Fig 4.12, which shows the profiles of electron temperature between 30-80 ns for

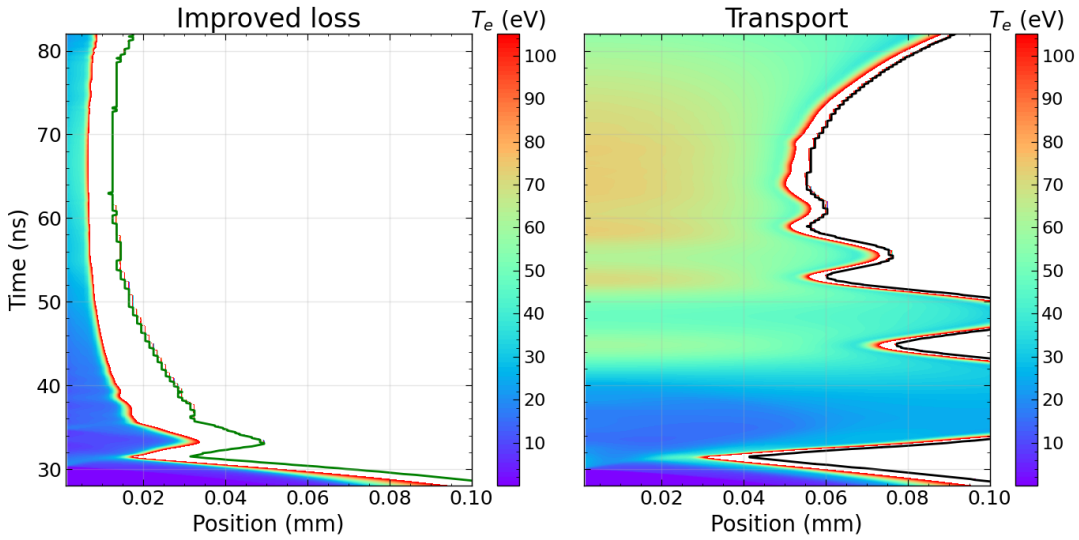


Figure 4.12.: Streak profiles showing electron temperature, zoomed in to $r \leq 100$ μm region, comparing late time (30-80 ns) dynamics of the wire using the improved radiation loss model (left) and the full transport case (right). Pseudo-colour range extends only to $T \leq$100 eV. Interface profiles are overlaid in green (left) and black (right)

the two models. Note, only states at $\leq$ 100 eV are shown on these plots for clarity, which mostly coincides with the core region. The loss case (left plot), while showing small local increases where the wire is repeatedly compressed, remains at a low value of $\leq$ 40 eV. Comparatively, the transport case (right plot) remains overall hotter, and reaches a peak value of 75 eV. In the same region the radiation temperature reaches a peak of 60 eV, comparable to the electron temperature and therefore suggests a state close to thermodynamic equilibrium.

The core densities in the transport case at the times shown are between 50-100 kg/m^3 here. The opacity spectrum for these conditions is shown in Fig 4.13, from the same atomic data used in simulations. The peak energy for a 75 eV Planckian spectrum is shown in the red line, and corresponds to $\kappa = 2 \times 10^5$ m^{-1}. An estimate for the photon mean free path at this energy is therefore $\lambda \sim 1/\kappa = 5$ μm, a factor of 10 smaller than the compressed radius, suggesting a considerable fraction of the radiation is trapped within the plasma volume.

4.3.4. Final comments

To summarise, it is apparent that using higher fidelity equation of state and transport models is required to predict the early time resistive phase, and as well as the rate of energy deposition into the core (expansion of core using FEOS and LMD models is 0.8 cm/μs most closely matches experimental rate of 0.72 cm/μs). Discrepancies against experiment include an overly aggressive rate of coronal expansion, which is $\sim$7 cm/μs in simulation and 2.2 cm/μs in experiment. Additionally, the group-resolved loss most accurately matches the early ablation and re-compression stages of the radiation

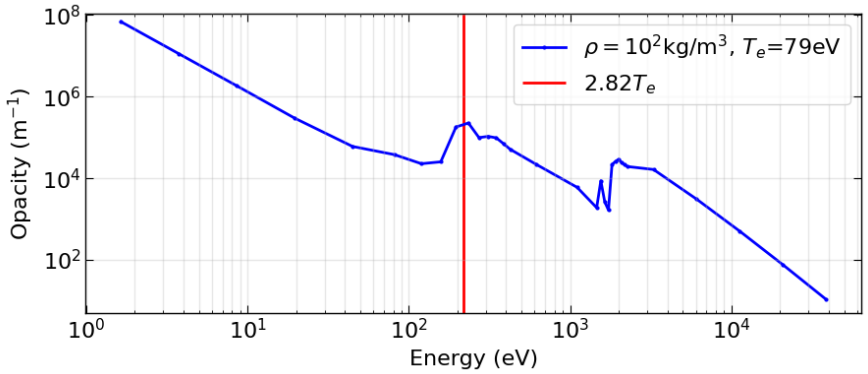


Figure 4.13.: Opacity spectrum for the core states in the radiation transport case at close to peak current.

transport reference case, although there are evident differences in the stagnated state between the different radiation models. The transition to an optically thick regime will not be captured by the loss model, resulting in an overestimate of the compressed state pressure. However, in the absence of a higher-fidelity yet similarly fast model this discrepancy will simply have to be accounted for in discussion. The model options used by default in all subsequent simulations (unless stated otherwise) are FEOS for equation of state, LMD for transport, and group-resolved loss for radiation.

4.4. 2D simulations

4.4.1. Set up

This section extends the above work to two dimensions. The simulations are performed in cylindrical geometry, with a radial and axial extent of $L_r \times L_z = 2 \times 2$ mm^2. Using mesh refinement the resolution on axis is set to $\Delta r = \Delta z = 1$ μm and 16 μm further away, corresponding to 5 levels of refinement. The excessively large domain in the $\hat{r}$-direction delays low density plasma hitting the magnetic boundary and generating numerical instabilities, allowing the simulation to run for longer at a small additional computational cost. A total of 2.5×10^5 grid cells were used in this splitting configuration, whereas an equivalent uniform grid at this minimum resolution would require 4×10^6 cells. In this instance the savings are not vital and a uniform grid could feasibly be used, however mesh refinement is used for comparison against 3D simulations where the savings are more significant (§4.5).

To verify that using mesh refinement does not drastically change the dynamics, a 1D comparison of a uniform mesh versus (only radially) refined mesh is performed. Fig 4.14 presents the equivalent streak plot in Fig 4.8 (for density rather than pressure) for a uniform grid simulation (left) versus a simulation run with mesh refinement ($\Delta r = 1-16$ μm, right), where the interface of the left is overlaid on the right. As expected the time and rate of the initial core expansion are correctly matched, as is the pressure and radius at the final compressed state. However, the coronal expansion deviates early in time

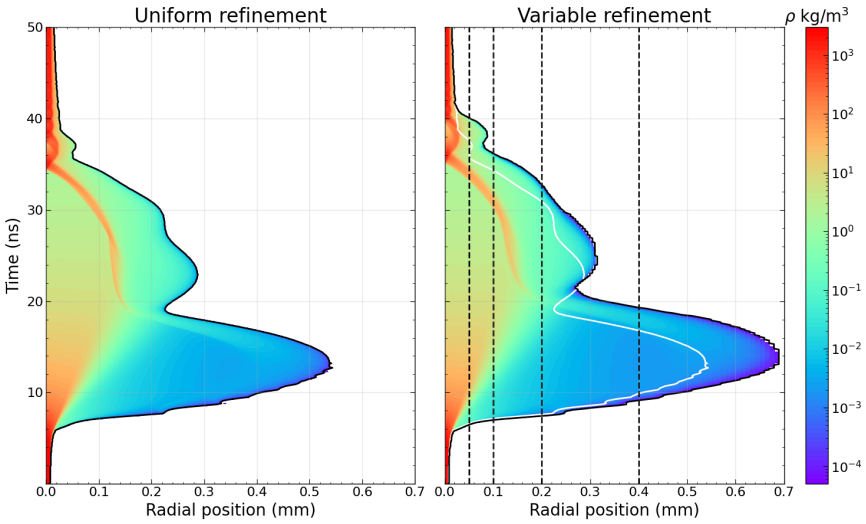


Figure 4.14.: Density streak plot comparison of uniform mesh (left) versus variable refinement mesh (right). Dashed vertical lines indicate resolution boundaries, i.e. between regions of $\Delta r = 1\mu$ m, 2μ m, 4μ m, 8μ m, and 16μ m. White line overlaid on right plot is the interface profile from the left plot.

when crossing through regions of the lower resolution. This results in further expansion and therefore a temporal offset in subsequent dynamics. Ultimately the high resolution regions closely match the uniform mesh case, barring some differences in the corona which do not strongly affect the compressed state.

Perturbations

The 1D simulations were initialised with solid density and ambient temperature states. In 2D, a small perturbation on the temperature profile was applied volumetrically through the wire to initiate instability formation. This took the form:

$$T_e = T_i = \frac{293}{11604} \times (0.99 + 0.02\epsilon)[\text{eV}], \tag{4.4}$$

where ϵ is a random number between 0 and 1. A perturbation in temperature was chosen as it would reflect in the conductivity and imitate inclusions in the material that seed ETI formation. These are shown in Fig 4.15(a), in the bottom two panels. The dynamics of the evolution were found to be sensitive to the amplitude of the perturbation. Initially it was applied to the temperature in Kelvin (292 ± 2 K), however the conversion to eV resulted in too small a variation in conductivity, giving a smooth profile. On the other hand, using a perturbation of ± 0.05 eV instead of ± 0.01 eV increased the maximum conductivity value from $4 \times 10^7 \Omega^{-1}$ m^{-1} to $4.5 \times 10^7 \Omega^{-1}$ m^{-1}, allowing quicker diffusion of the field into the wire and initiating earlier expansion.

4.4.2. Dynamics

The wire evolution mostly resembles the 1D simulations for the early time, axially uniform phase of growth. However at late time 2D, non-linear $m = 0$ growth results in substantial deviations. Figs 4.15(a)-(c) shows a magnified view the early time linear phase, and Figs 4.16(a)-(c) show the full view at a later time, both are considered individually.

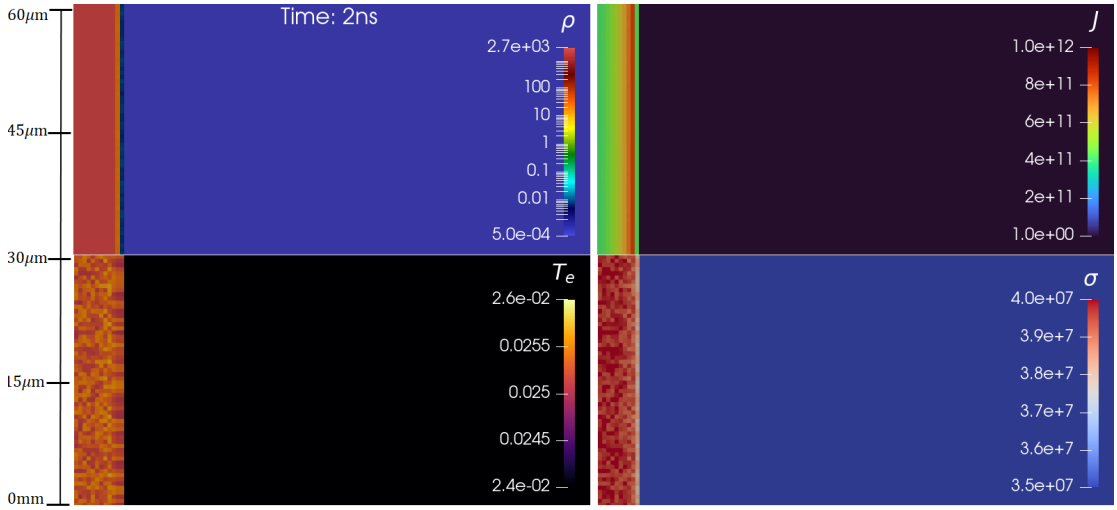
Early-time

Fig 4.15(a) shows the simulation 2 ns after start. The current is diffusively moving through the wire which slightly heats its surface, although the rate of Ohmic heating is low as the wire is still in a conductive regime. The sensitivity of the equation of state means that even a slight temperature increase raises the pressure from near zero to 10^6 Pa. The surface begins to rarefy, and convectively carries its temperature perturbation along with the radial velocity. By 8 ns (Fig 4.15(b)) the wire has expanded, and is undergoing a transition towards a core-corona structure. This is evidenced by the inversion of conductivity in the radial direction, which causes the current to switch from core (as in Fig 4.15(a)) to corona (as in Fig 4.15(c)). While the initial perturbations appear smoothed out in these plots, this is due to the visible range of perturbations being much smaller than the full range shown. The next output at 9 ns (Fig 4.16(c)) shows a rapid expansion to twice the previous radius, with a hot, current-carrying corona fully formed. Also visible is the formation of short wavelength modes ($\lambda \sim 20$ μm) along the interface, despite a relatively smooth profile further inwards. This is discussed further in §4.4.3.

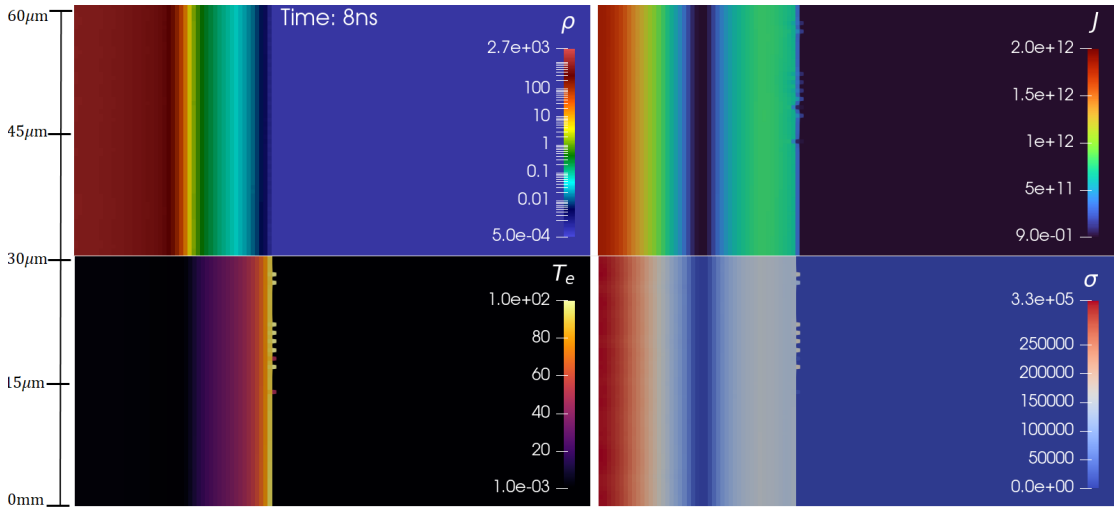
Late-time

Figs 4.16(a)-(c) shows the dynamics after Fig 4.15(c) across a larger temporal range and larger spatial extent. Fig 4.16(a) at 10 ns gives a sense of scale of Fig 4.15(c), when the wavelength of perturbations are of order $ka \sim 1$, after which point this relation does not hold. For the next 15 ns, the wire continues expanding, maintaining a smooth interior but unstable leading edge. Rather than expand until 14 ns then contract (as in 1D simulations), the expansion continues until closer to 18 ns, after which certain regions decelerate and others continue to expand. This is due to the $m = 0$ instability, where neck regions of slightly smaller radii undergo a feedback loop of contracting under increasing magnetic pressure, and bulges exceed a critical radius above which they do not get magnetically confined. By 25 ns, the contraction has caused a local increase of current density on the leading edge, although it is enhanced in the necks and reduced in the bulge, forming a similar current path as described in Fig 4.3. As in 1D, the current-carrying plasma gets increasingly compressed against the core between 30-40 ns. The diffusion of current into the core generates higher temperature and pressures which resists the compression. The reflected shock off the axis causes non-linear Richtmyer-Meshkov mixing of the surface, and generates multi-wavelength mode perturbations.

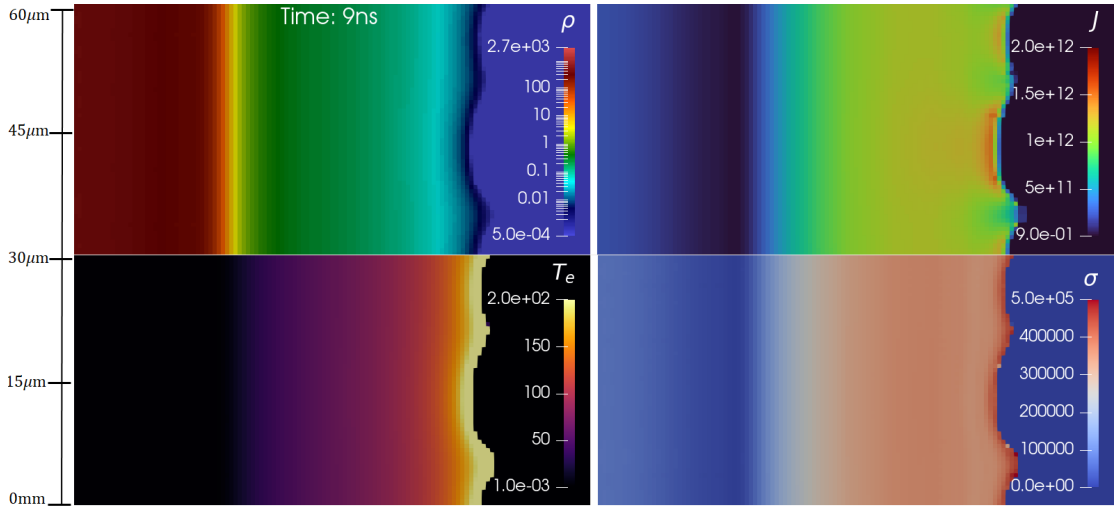
In regions where the compression is not sufficiently decelerated, the wire forms regions of intense pinching at ∼40 ns. This is enabled by the extra degree of freedom of motion in the axial direction, which depletes mass from the neck until finally it bifurcates to form separate regions. Peak compression therefore occurs earlier at 41 ns rather than at peak current as in the 1D simulations. A few nanoseconds afterwards low density surrounding plasma fills the gaps where the neck was, however is at too large a radius to experience a strong $\mathbf{j} \times \mathbf{B}$ force. By 60 ns the behaviour is highly non-linear and low density plasma has filled most of the simulation domain so no further compression occurs. The simulation is terminated before reaching peak current.



(a) $t = 2$ns

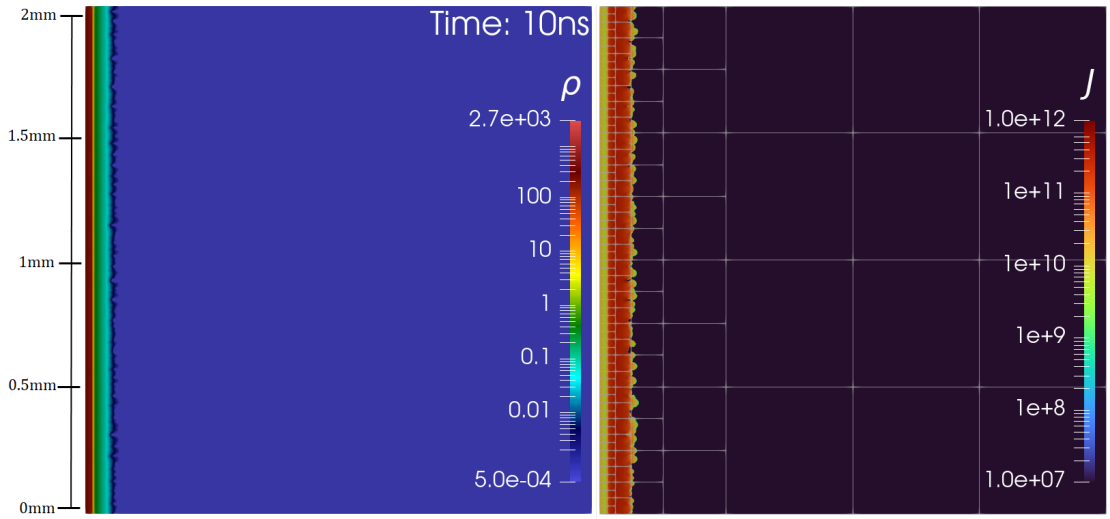


(b) $t = 8$ns

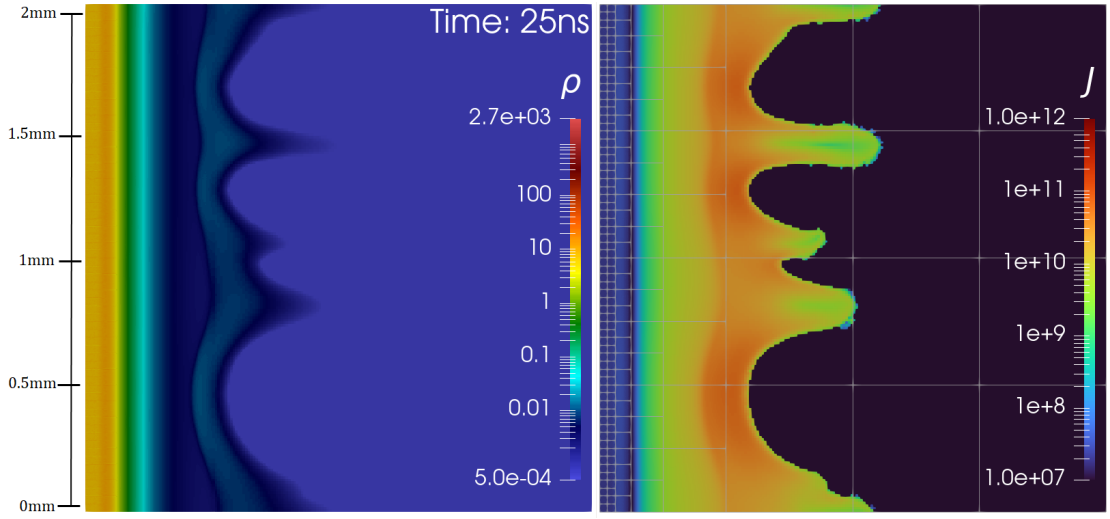


(c) $t = 9$ns

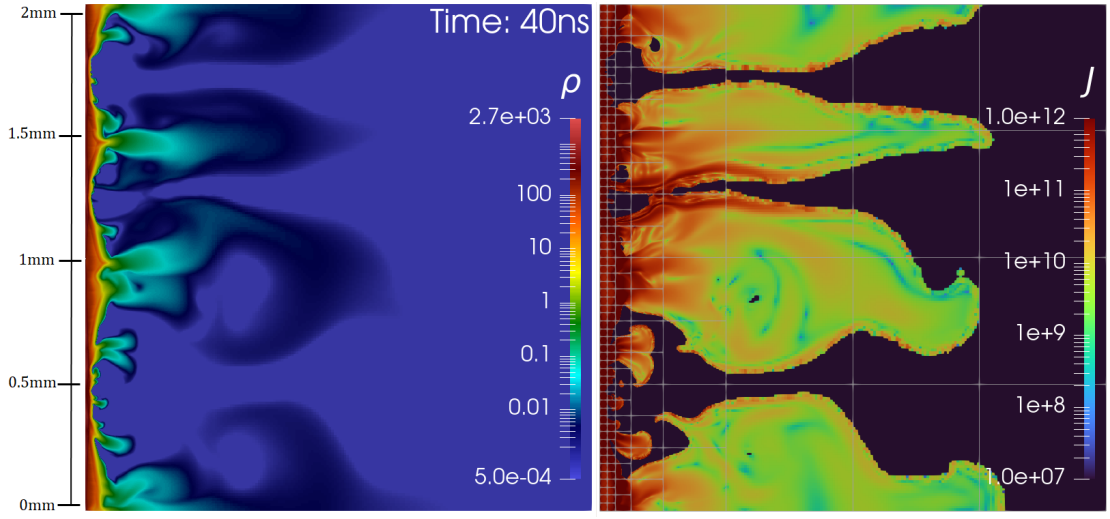
Figure 4.15.: Early time evolution of 2D simulations of 15 μm Al wire dynamics, showing close-up profiles at three different times. Each image shows the log mass density (top left), log current density (top right), electron temperature (bottom left), and electrical conductivity (bottom right).



(a) $t = 10$ns



(b) $t = 25$ns



(c) $t = 40$ns

Figure 4.16.: Late time evolution of 2D simulations of 15 μm Al wire dynamics, showing full domain at three different times. Each image shows the log of mass density (left) and current density (right).

4.4.3. Wavelength analysis

Fig 4.16 shows that overall the coronal plasma perturbations increase in wavelength over time, similar to the series of Schlieren images in Fig 4.5. The modulations in the experimental images at early time are less discernible, however between 36-48 ns the images show wavelengths between $500 - 800$ μm. This is comparable to the ~ 660 μm wavelength observed in simulations at 40 ns (Fig 4.16(c)).

The uncertainty regarding the seed of the $m = 0$ instability can hypothetically be addressed with the high spatial resolution accessible using mesh refinement, however these simulations were not run with sufficient temporal resolution to capture the early growth. Instead, some observations are made about the existing output. Fig 4.17 shows the wire at 7 ns and 9 ns, immediately before and after the current switches into the corona as demonstrated in Fig 4.16. Shown from left to right are the radial current, axial velocity, gradient of electron temperature in the axial direction, and log density. Colour bar ranges are set to more clearly visualise the positive/negative variation.

At 7 ns, the first three show striations with axial wavelengths at approximately $\lambda \sim 5$ μm, i.e. of order 5 cells across. This is still considered resolution limited, therefore future investigations should to study core structures should be run using higher resolutions.

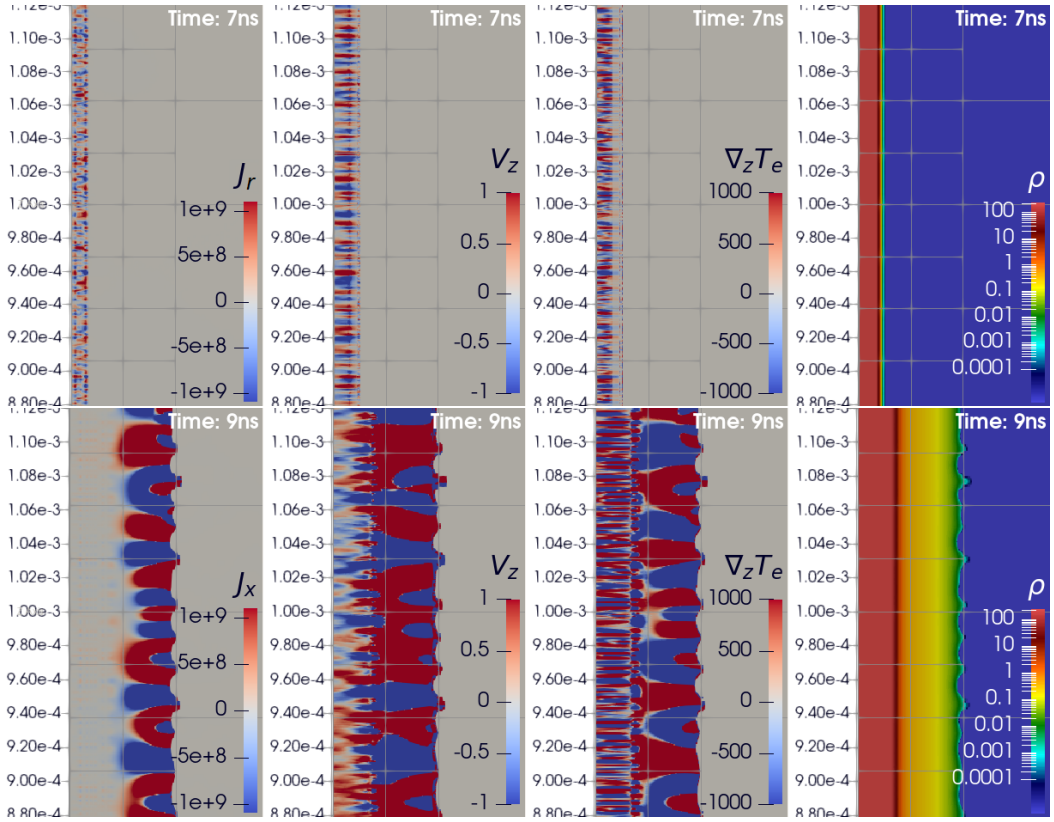


Figure 4.17.: Axial variation of radial component of current density (J_r), axial velocity (v_x), axial electron temperature gradient ($\nabla_z T_e$), and log density (ρ) at 7 ns (top row) and 9 ns (bottom row). Axial extent of the magnified view is 240 μm.

The banded structures in J_r are less radially correlated than in v_z and $\nabla_z T_e$, suggestive of a left-to-right weaving motion that could arise from the ETI. The axial velocity is anti-correlated to the temperature gradient, suggesting a hydrodynamic expansion from this variation. Both the v_z and J_r components are 2-3 orders of magnitude smaller than the other directions (v_r and J_z respectively), so this effect is a small perturbation to the bulk motion. At 9 ns, the coronal plasma has considerably different axial variation profiles to the core. The core retains its small wavelength structure in electron temperature gradient, while the bands in axial velocity appear to merge and increase in wavelength. These are still of a shorter wavelength than the coronal plasma variation, which shows a modulation with a characteristic scale length $\lambda \sim 20-30$ μm across all variables. The temporal resolution is insufficient to capture how these longer wavelength modes form from the short wavelength variations in the core, however the perturbations in the two regions remain distinct, until later times when the corona impinges back on the core and imprints longer wavelength modes.

The core and corona follow a complicated evolution from this time on, and it becomes intractable to consider each variable as a function of time in 2D. A typical technique for wavelength analysis is to plot the mass per unit length as a function of axial position [101]. However, radially integrating through the domain on an unstructured cylindrical mesh requires a bespoke tool that has not yet been developed. Instead, the surface contour at $\rho = 10^{-4}$ kg/m^3, which corresponds to the plasma-vacuum interface, and

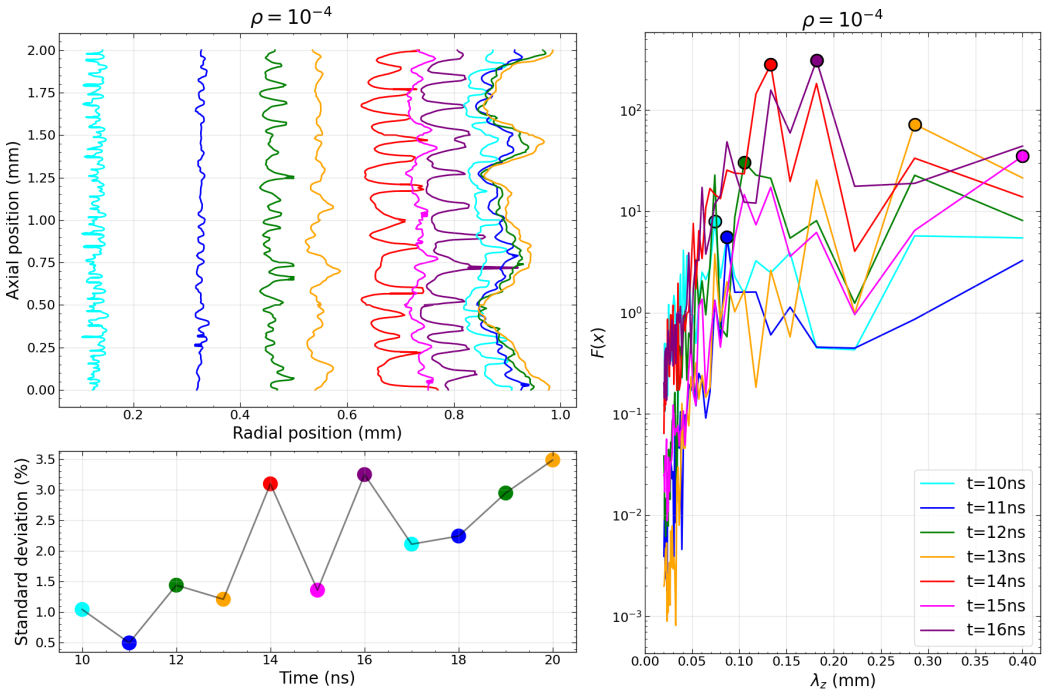


Figure 4.18.: *Top left:* spatial profiles of $\rho = 10^{-4}$ kg/m^3 contour as a function of time. *Bottom left:* standard deviation of profiles above as a function of time. *Right:* Fourier transforms of spatial profiles, showing the highest amplitude wavelength at each time in coloured dots.

$10 - 20$ ns is shown in the top left plot of Fig 4.18.

As the corona expands it undergoes changes in wavelength and amplitude. The small-scale perturbations at 10 ns have $\lambda \sim 20 - 30$ μm but are otherwise uniform in radial position. By 12 ns, the wavelength is doubled but still centred on a uniform equilibrium value. The profile at 16 ns has evidence of small scale features but forms long wavelength modes, around 2-3 oscillations within the 2 mm axial extent. The later-time set of cyan, blue, green, and yellow lines ($t = 17 - 20$ ns) show the corona decelerating and beginning to contract. By this time the profile is primarily a long wavelength mode with superimposed fine-scale oscillations.

While the spatial profiles mostly show an increase in the wavelength over time, this is not reflected in the spectral decomposition. Fourier transforms of the profiles from $t = 10 - 16$ ns are shown on the panel on the right, where dominant wavelengths (highlighted in coloured dots) do not follow the observed trend. The reasons for this are twofold; firstly the Fourier signals in the small wavelength range show significant noise, and secondly the long wavelength modes are artificially weighted as this corresponds to the domain size. Furthermore, resolution effects are also challenging to identify as some profiles span regions of variable resolution. It is expected that performing the same analysis on spatially integrated, linear-mass data will give clearer signals, which can then be used to extract parameters such as ka as a function of time.

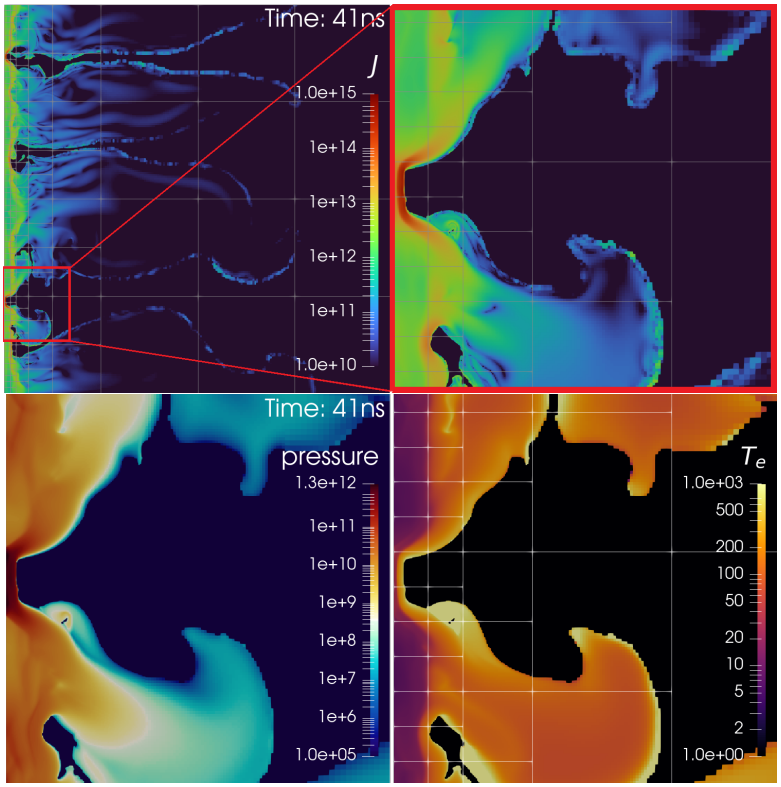


Figure 4.19.: Compressed state of a single pinch point along the fibre at $t = 41$ ns. *Top:* logarithm of current density showing the full wire and a close-up profile (550×550 μm^2 region). *Bottom:* pressure and electron temperature profiles of the magnified region.

4.4.4. Hotspot

The state of the hotspot is shown in Fig 4.19 at the time of peak compression. Multiple pinch points are formed along the wire, only a single one is considered here. Before bifurcating, the pressure at the pinch point reaches a maximum of > 1 TPa with solid density conditions over a 2 μm radius hotspot. Based on the analysis in §4.3.3, this figure may be an overestimate of the value reached using the radiation transport model. This compression is exerted on the high density core by a current carrying region, that experiences a strong magnetic field of order 10^3 T. Temperatures vary between 20-600 eV across both regions, however the region with the highest radiative cooling rate is at $r = 5$ μm, at a hundreth of solid density and 80 eV. This rate is adjusted by a probability of escape factor, and therefore marks the point at which the density scaling of the radiative loss is decreased due to optical depth effects.

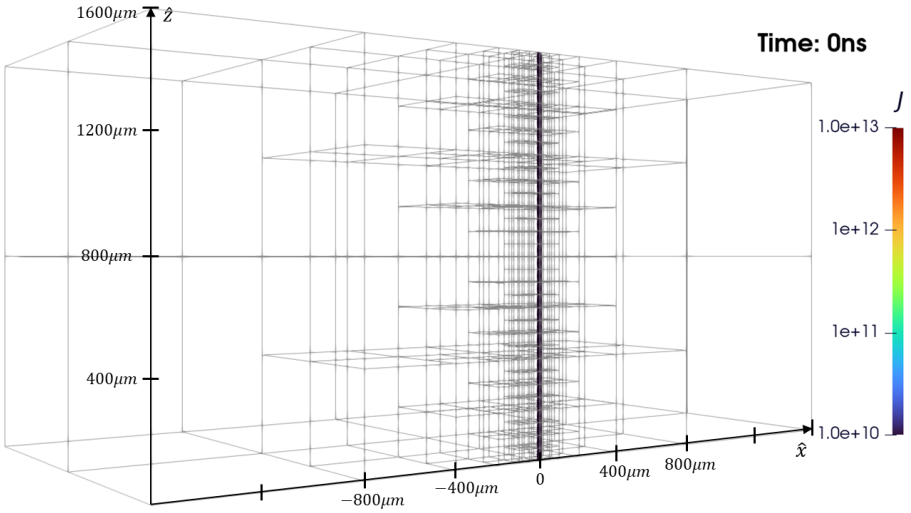
4.5. 3D simulations

4.5.1. Setup

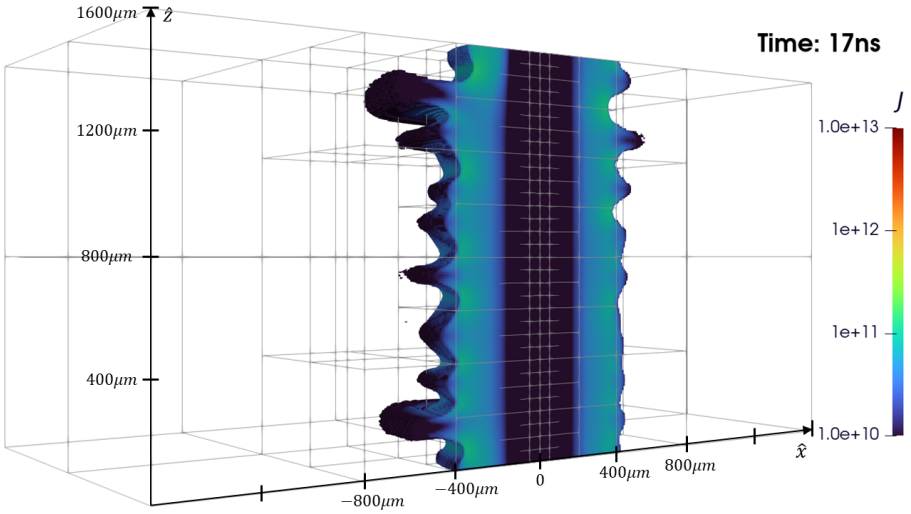
This section describes the 3D Cartesian simulations of the above system across a domain of $L_x \times L_y \times L_z = 3.2 \times 3.2 \times 1.6$ mm^3, again using mesh refinement with the same levels of refinement; $\Delta x = \Delta y = \Delta z = 1$ μm along the axis and 16μm away from the centre. In this refinement configuration a total of 872 processors were used, each with $50 \times 50 \times 50$ cells. Running an equivalent uniform resolution simulation at 1μm with the same number of cells per processor would require $64 \times 64 \times 32 = 131,072$ processors, giving savings of over 150 times.

4.5.2. Dynamics

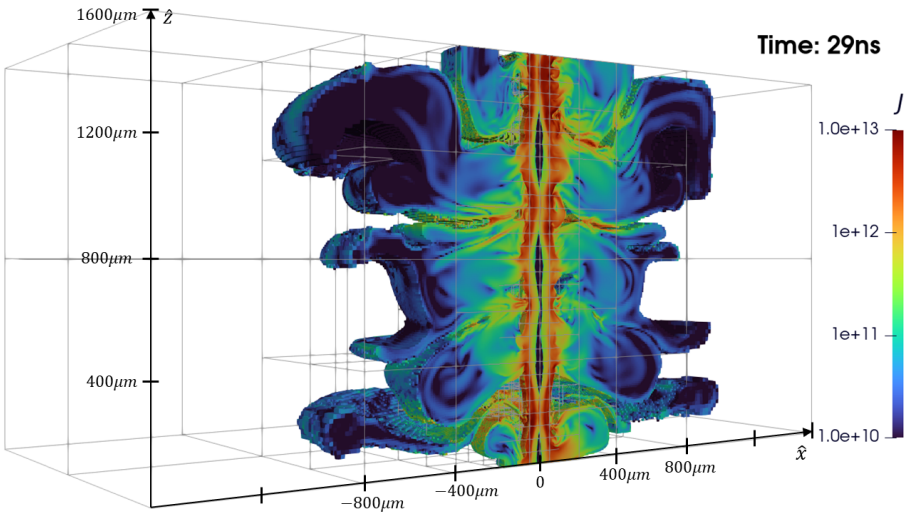
The evolution of the system is shown in Fig 4.20, showing the initial conditions and mesh splitting in the first subplot. Note, these simulations used a slightly larger temperature perturbation, resulting in a larger conductivity variation and therefore a stronger initial rate of Ohmic heating. As a result, the initial rate of expansion is faster than in the 2D case and all subsequent dynamics occur at least 2 ns earlier. This expansion phase reaches a peak at 500 μm rather than 800 μm as in 2D, and shows only little azimuthal asymmetry at this point. The compression of the corona onto the core, and the subsequent onset of non-linear shock-driven mixing, occurs at least 5-10 ns earlier than the 2D case at 25 ns. By 29 ns, the current has accumulated to just outside the core, and has even penetrated the core in certain regions. The low density region has become strongly turbulent, and bulges of the $m = 0$ asymmetry are rapidly ejected outwards. These regions, however, carry little current relative to the central column. Despite the non-linearity the pinch is still remarkably symmetric in the azimuthal direction. The asymmetry arises afterwards, as described below.



(a) $t = 0$ns



(b) $t = 17$ns



(c) $t = 29$ns

Figure 4.20.: Evolution of single wire simulated in 3D at three different times, shown in terms of log current density. Also overlaid is the mesh splitting, shown in grey lines.

4.5.3. Azimuthal symmetry

Fig 4.21 shows the streamlines of the magnetic field across the central column of the pinch at different points in its evolution. The images here are intended to show a qualitative overview of the behaviour rather than identifying particular wavenumbers etc. What is evident, however, is that an $m = 1$ instability begins to form only at later times once the current is compressed to small radii. The lower half appears to go unstable earlier than the top half at 30 ns, however by 35 ns the entire column is unstable. The core is broken through the $m = 1$ characteristic twisting motion, and forms distinct high density regions (rightmost panel of Fig 4.21), reminiscent of the later time Schlieren images in Fig 4.5. This motion further disrupts hotspot formation, except in one high density region at the lower end, as discussed below.

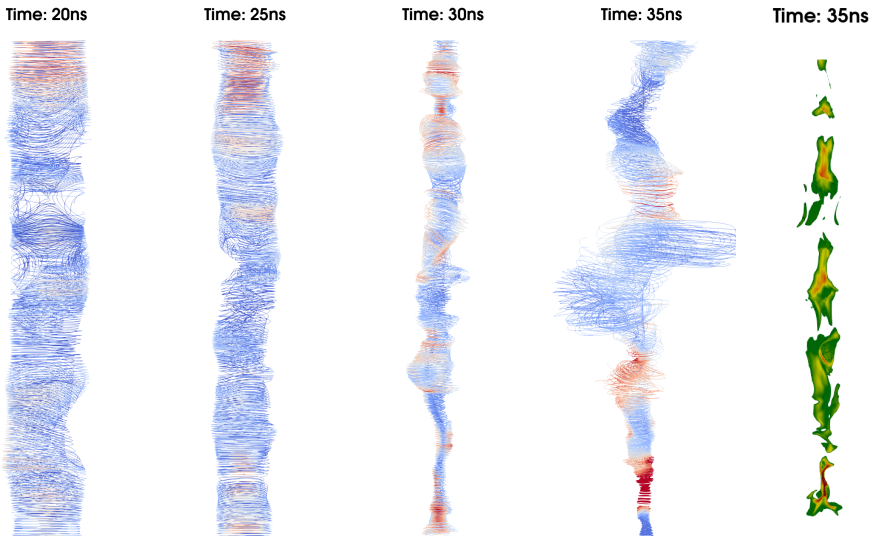


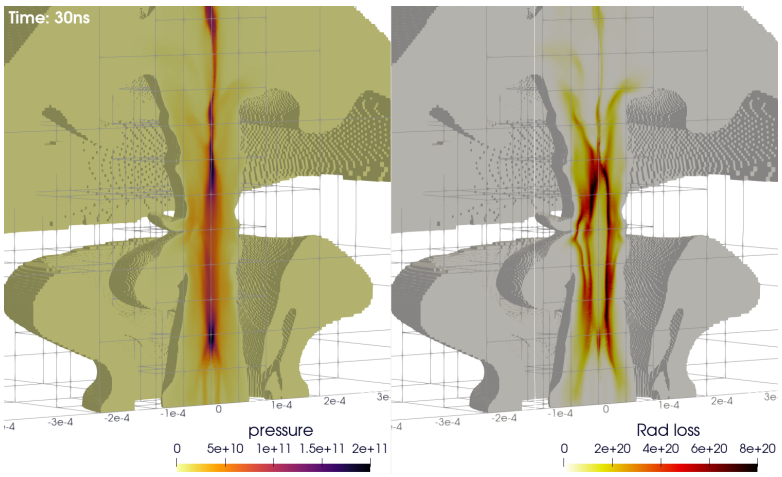
Figure 4.21.: Evolution of $m = 1$ instability across the full length of the Z-pinch, shown in terms of streamlines of the magnetic field. Colour scale indicates magnitude of magnetic field. Initial expansion is highly uniform, late time compression of corona into the core forces. *Rightmost image*: slice of density profile thresholded at $\rho > 10^2$ kg/m^3, showing high density core broken up into separate islands.

4.5.4. Hotspot(s)

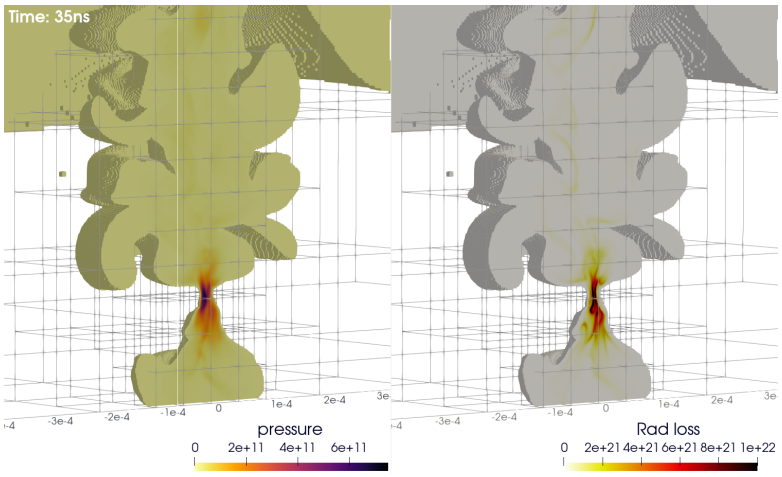
Between 30-35 ns, the wire undergoes compression through two difference mechanisms: first due to compression from an $m = 0$ instability, and the second via an $m = 1$ kink instability. The first forms at 30 ns near the lowest neck along the wire ($z \approx 200$ μm in Fig 4.20(c)), shown from a magnified view in Fig 4.22(a) in terms of pressure and radiation loss. The left image shows two distinct regions of high pressure at higher and lower positions along the axis, which have values 1.9×10^{11} Pa and 2.1×10^{11} Pa respectively. These are an order of magnitude lower than the hotspot described in §4.19 in 2D simulations (1.3×10^{12} Pa), this difference is discussed below. The size of the higher pressure (further down the axis) hotspot is 10 μm across and $\sim$ 15 μm long

for the $> 1.5 \times 10^{11}$ Pa region, a factor of at least 5 greater than the 2D simulations. The highest loss rate as estimated by the improved loss model does not come from the hotspot itself, but rather the surrounding regions, see for example the bright emission along the edges of the core in Fig 4.22(a).

At 31 ns, the high pressure core appears to have rebounded rather than bifurcate. The pressure relaxes and the currrent-carrying region slightly moves apart. However, subsequent timesteps show a re-pinching. By 35 ns, the current has switched entirely into the core, creating twisting motions from the $m = 1$ instability that moves the flow in non-radial directions. At one point in its motion this behaviour forces a compression of the core to an 8 μm region that reaches even higher pressures than the hotspot at 30 ns, at 7.5×10^{11} Pa. However, this is disrupted shortly afterwards by the continued stretching along the axis. No further compressed states form, as by this time majority of the core has disintegrated and the corona has expanded to hit the domain boundaries where the magnetic field boundary conditions are set, causing the simulation to terminate.



(a) $t = 30$ns



(b) $t = 35$ns

Figure 4.22.: Magnified view of hotspot formation in 3D simulations of a 15 μm Al wire at two different times, in terms of pressure (left) and radiative loss (right).

4.6. Concluding remarks

This chapter presents a numerical study of a single 15 μm Al wire exploding in vacuum using the IMP generator which delivers a 150 kA peak current, 65 ns rise time pulse. These experiments were performed several years ago at Imperial College, at which time they were simulated with the best available models. The availability of improved EoS and transport models provide impetus to revisit these systems. Additionally, the complex evolution of this ostensibly simple system has not previously been studied in 3D due to the intractable resolution requirements, an obstacle which the mesh refinement capability described in the previous chapter overcomes. Therefore, a systematic study of this system from one, to two, and finally three dimensions was conducted.

The 1D behaviour shows an early fast expansion of the corona until 14 ns to 400 μm, after which cycles of smaller compression and re-expansions occur. The corona is compressed to the core during peak current, after which point it is free to expand. In the radiation transport case, the stagnated state reaches a minimum of $r_{min} = 55$ μm, while in the loss cases they reach 13 μm. The difference of pressure in the stagnated state is scales roughly as the inverse of this ratio squared, 0.23 TPa versus 2.7 TPa respectively. While this is a significant discrepancy, running with radiation transport at these conditions and resolutions is challenging and will considered as future work. The 2D simulations introduce an additional degree of freedom to the plasma motion. This results in significant mass loss from the hotspot region, resulting in 1.3 TPa however over a $r = 2$ μm region. While it is clear that strong compression occurs, bifurcation of the hotspot appears to terminate any further compression to a radiatively 'collapsed' state. Including an additional dimension then introduces an additional termination mechanism, the $m = 1$ instability, which drops the peak pressure by an extra factor of 2. The 2D and 3D simulations were run with the loss model, therefore possibly still overestimate the compression reached. Conversely, they were run with relatively low temporal resolution compared to the pinch dynamics, and peak compression may not be observed with the available output. Regardless, this work has demonstrated the ability to comprehensively model a radiatively cooled system at high resolution. This capability is tested on the X-pinch system next, for which only 3D modelling is appropriate to capture dynamics.

5. X-pinches

The multiple hotspots formed in single wire experiments described in the previous chapter form randomly along wire length, but can be directed to a single point using X-pinch devices. X-pinches were first fielded on the BIN generator at the Lebedev Physical Institute in 1982 [104], and have been widely used since. The typical configuration used is formed by taking a standard wire array of two (or more) wires, and twisting the top or bottom electrode by 180° such that the wires are angled and cross in the centre. The current is thus focused at this crossing point, creating intense plasma conditions even with relatively modest generators. As in standard wire arrays, the individual wire behaviour is significantly different to the single wire dynamics previously considered, due to the resultant global magnetic field topology that arises from the load configuration.

The work in this chapter explores the hotspot formation in X-pinch devices, which have long been speculated to undergo radiative collapse. Two sets of experiments and corresponding configurations are considered here. The first is a 2-wire X-pinch in the discrete wire configuration described above, performed on XP at Cornell University, a 400 kA generator. The second is a larger mass, solid machined X-pinch, performed on SATURN, a MA-class generator. The former has been extensively studied experimentally and is well characterised, however is challenging to simulate correctly due to stringent resolution requirements. The latter is difficult to diagnose experimentally because of the extreme currents involved, however has an intrinsic symmetry that makes it easier to simulate. Both undergo strong radial convergence to reach extreme conditions at a fixed point, making them ideal systems to study using a static mesh refinement capability.

This chapter is organised as follows. First, an overview of the wire X-pinch configuration is provided in §5.1, as this is the most commonly fielded configuration. This includes a review of the typical dynamics observed, plasma conditions reached, and common applications of the devices. §5.2 and §5.3 then describe the 400 kA and MA-class experiments respectively. In both cases an overview of the experimental results and diagnostics is first provided, followed by a description of the simulation set-up, and finally discussion of results.

5.1. Overview

The X-pinch has been studied in a variety of configurations and on a large range of experimental facilities. The two most common configurations are the wire X-pinch (WXP, described above) and the 'hybrid' X-pinch (HXP). These, as well as laser-cut

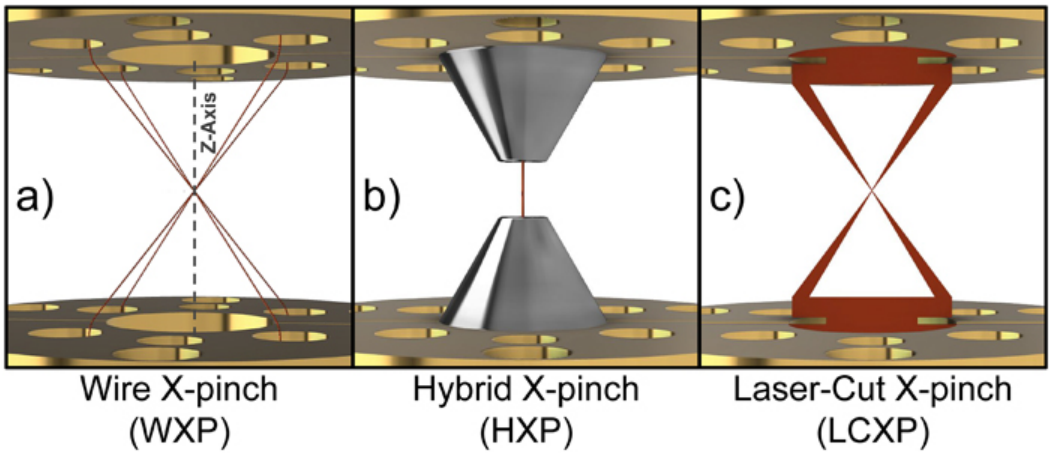


Figure 5.1.: Illustration of three different X-pinch configurations. The WXP (leftmost) is the focus of subsequent sections. Reproduced from Ref [19].

X-pinches (LCXP), are illustrated in Fig 5.1. The HXP is a rotationally symmetric extension of the wire pinch, and uses two oppositely facing conical electrodes connected via a single wire to achieve a similar focussing effect. The LCXP, as the name suggests, uses a laser-cut thin foil to create tapered legs and therefore reducing mass per unit length towards the centre. Of the three, the WXP has been most extensively studied in literature, and is therefore the focus of this section and §5.2. A variant of the HXP is considered in §5.3.

5.1.1. Dynamics

WXPs have been extensively studied on a range of experimental facilities across the world, and their dynamics have been well characterised, often using secondary X-pinches to radiograph a target one [21]. The evolution observed from these experiments typically occurs in four distinct stages, described as follows using Fig 5.2 as a reference.

1. As described in §4.1 for single wires, individual wires in the WXP do not undergo uniform volumetric heating and expansion. Instead, a core-corona structure is similarly initiated from early time. The coronal plasma is swept inwards towards the axis to form the typical observed axial jet structure, highlighted in the left-most panel of Fig 5.2. This is due to a combination of two distinct effects. Firstly, the global magnetic field develops such that the resultant $\mathbf{j} \times \mathbf{B}$ force is oriented perpendicular to the wire cores, and directed towards the axis, as shown in dashed grey arrows in the diagram. Secondly, a magnetic null is formed on the axis on either side of the crossing point where oppositely directed fields from the wire cancel. Therefore when the crossing point is Ohmically heated, there is no confining magnetic field to oppose the pressure gradient $\partial P/\partial z$ (blue arrows in the diagram), and the plasma is free to expand in the axial direction.

2. The magnetic field reaches a maximum at the crossing point, i.e. the point of minimum radius, and the Lorentz force is directed entirely radially, resulting in

enhanced compression. Eventually this forms a 'micro Z-pinch' structure, typically spanning 100s of μm in the axial direction, which evolves separately to the 'legs' of the X-pinch (second panel).

3. Upon further compression, the micro Z-pinch becomes $m = 0$ unstable and forms a series of necks along its length. Around this time intense bursts of soft X-rays ($\lesssim$ 1 keV) from highly compact hotspots are observed (third panel in Fig 5.2). The micro Z-pinch then bifurcates where the necks were formed, leading to plasma islands where the previously axially uniform structure was. Following the x-ray burst the hotspots explode radially and launch spherical shocks into the legs.

4. The shocks evacuate mass from the centre of the micro Z-pinch, forming a gap or 'mini-diode' between the legs on either side (last panel in Fig 5.2). Electrons are accelerated across the gap between the 'electrode' regions of higher plasma density (where the legs are still intact), and emit hard x-ray ($\geq$10 keV) bremsstrahlung radiation. The distance between 'anode' and 'cathode' is of the order of the length of the micro Z-pinch, corresponding also to the size of the observed emission region.

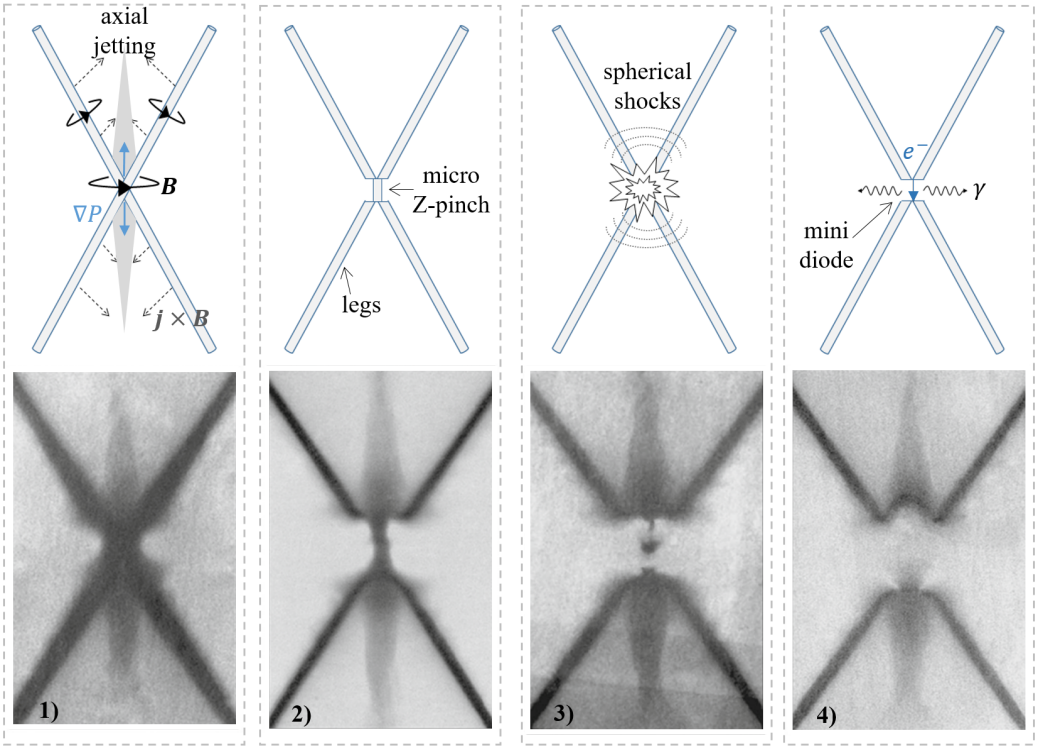


Figure 5.2.: Different stages of X-pinch evolution, shown diagrammatically along the top row with corresponding experimental radiographs along the bottom row. From left to right: 1) formation of core-corona structure along wires, where coronal material is swept towards axis to form an axial jet 2) radial compression of neck from Lorentz force to form a central micro Z-pinch 3) soft x-ray burst from hot spot(s) 4) hard x-ray emission from electrons accelerating across gap. Adapted from Ref [20].

5.1.2. Hotspot conditions

A comprehensive review of the spatial, temporal, and spectral properties of the X-pinch hotspot are provided in Ref [105], certain key observations are repeated here. These properties are sensitive to a range of parameters such as generator current, number of wires used, wire material, etc. However, spectroscopic analysis performed across several experiments has repeatedly demonstrated electron temperatures of $\gtrsim 1$ keV and number densities between $n_e \sim 10^{22} - 10^{23}$ cm^{-3}, where the emission originates from sources a few μm in size as estimated from self-emission images. Spatially integrated, temporally resolved characteristics of the x-ray emission are measured using photoconducting diodes (PCDs). Typical yields from these measurements on the XP (400 kA facility) are 0.4-1 J at energies > 1 keV, with pulse widths of the order of 100 ps. Both spatial and temporal properties reach the resolution limits of the respective diagnostics, suggesting even higher yields from smaller sources may occur in experiments.

5.1.3. Applications

The behaviour described above indicates a highly complex system, with steady-state conditions in the axial jetting, and simultaneous extreme, transient states formed in the crossing point. The axial jets have not been studied as widely as the hotspot conditions, however these are relevant to astrophysical systems such as outflows of active galactic nuclei (AGNs) [106]. Similar to the jetting from conical wire arrays, which are typically used for laboratory astrophysics experiments, the axial jets from X-pinches have high Mach numbers and undergo strong radiative cooling [107]. While these features exhibit several interesting properties, this work will focus on the central hotspots.

The compact sizes and strong emission make X-pinch hotspots useful for several other applications. These can be used to study conditions of astrophysical objects, for example estimated pressures in the centre of the hotspot approach the 100s of Mbar reached in white dwarf stars due to electron degeneracy pressure [108]. The ability to form highly compressed states means X-pinches have also been used for neutron production. Deutrated fibres have been fielded on a 380 kA generator at Cornell University [109], and separately on a 1.7 MA generator at the Kurchatov Institute [110]. Neutrons yields of up to 10^{10} have been demonstrated on the latter, however the mechanism by which these were obtained is ambiguous, and more likely to be beam-target as indicated by the observed anisotropy on neutron detectors. Finally, X-pinches are most widely used as source for point-projection radiography due to the high spatial and temporal resolution they offer. The 'pinch-time' can be tuned by adjusting the crossing-point mass for a given current through the load, allowing them to be synchronised against other experiments. They have been fielded on the return posts of wire array experiments [111], used as a secondary diagnostic X-pinch to study a target one [20], and also inductively driven on the Z generator to 200 kA [112]. The experiments described in the next section are an example of the second case, using up to three X-pinches simultaneously to obtain multiple radiographs of each X-pinch at different times in their evolution.

5.2. 200kA fibre X-pinches

This section investigates results from Ref [21]. Radiography data was obtained of wire X-pinches at several points in their evolution, making it a useful dataset for comparison against simulation. The driver in these experiments was the XP generator at Cornell University, delivering up to 460 kA in a 40 ns rise time, corresponding to an average $dI/dt \sim 10^{13}$ A/s. An overview of the experiment and radiography results are provided here for later comparison to simulations.

5.2.1. Experimental overview

Set up

In this campaign, a total of three X-pinches per experiment were driven using the configuration illustrated in Fig 5.3. Here, two central X-pinches (labelled XP1 and XP2) were driven in parallel, and therefore assumed to receive half the machine current each. The return current can was divided into three sections, one of which was connected to the cathode via a third X-pinch (labelled XP3). XP3 therefore received a third of the current, and was massed higher than the other two X-pinches. Using this configuration, 3 sets of radiographs could be obtained from the main burst of each X-pinch. For example (assuming XP1 is massed lower than XP2):

- XP1 (source) would radiograph XP2 (target) before XP2 pinched,
- XP2 (source) would radiograph XP3 (target) before XP3 pinched,
- and XP3 (source) would radiograph either XP1 or XP2 (target) after they pinched.

By obtaining multiple radiographs per experiment, an extensive time series was obtained that demonstrates the evolution of X-pinches with a time resolution between $\sim 0.5 - 5$ ns, shown in Fig 5.4.

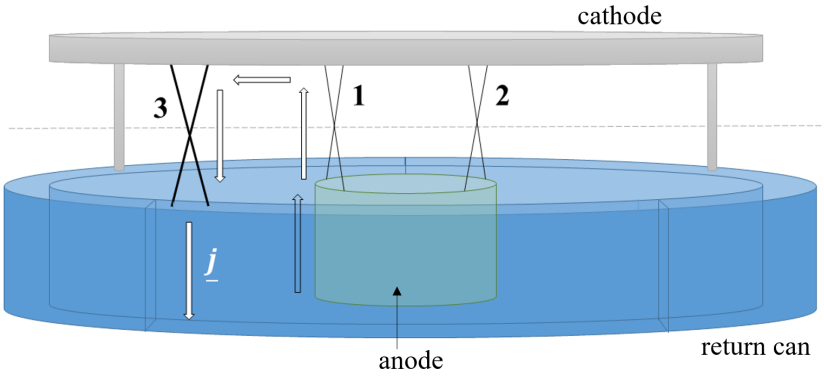


Figure 5.3.: Diagram of experimental set-up in Ref [21] showing two central X-pinches (labelled 1 and 2) and a third X-pinch (labelled 3). Direction of current from anode to cathode is shown in white arrows.

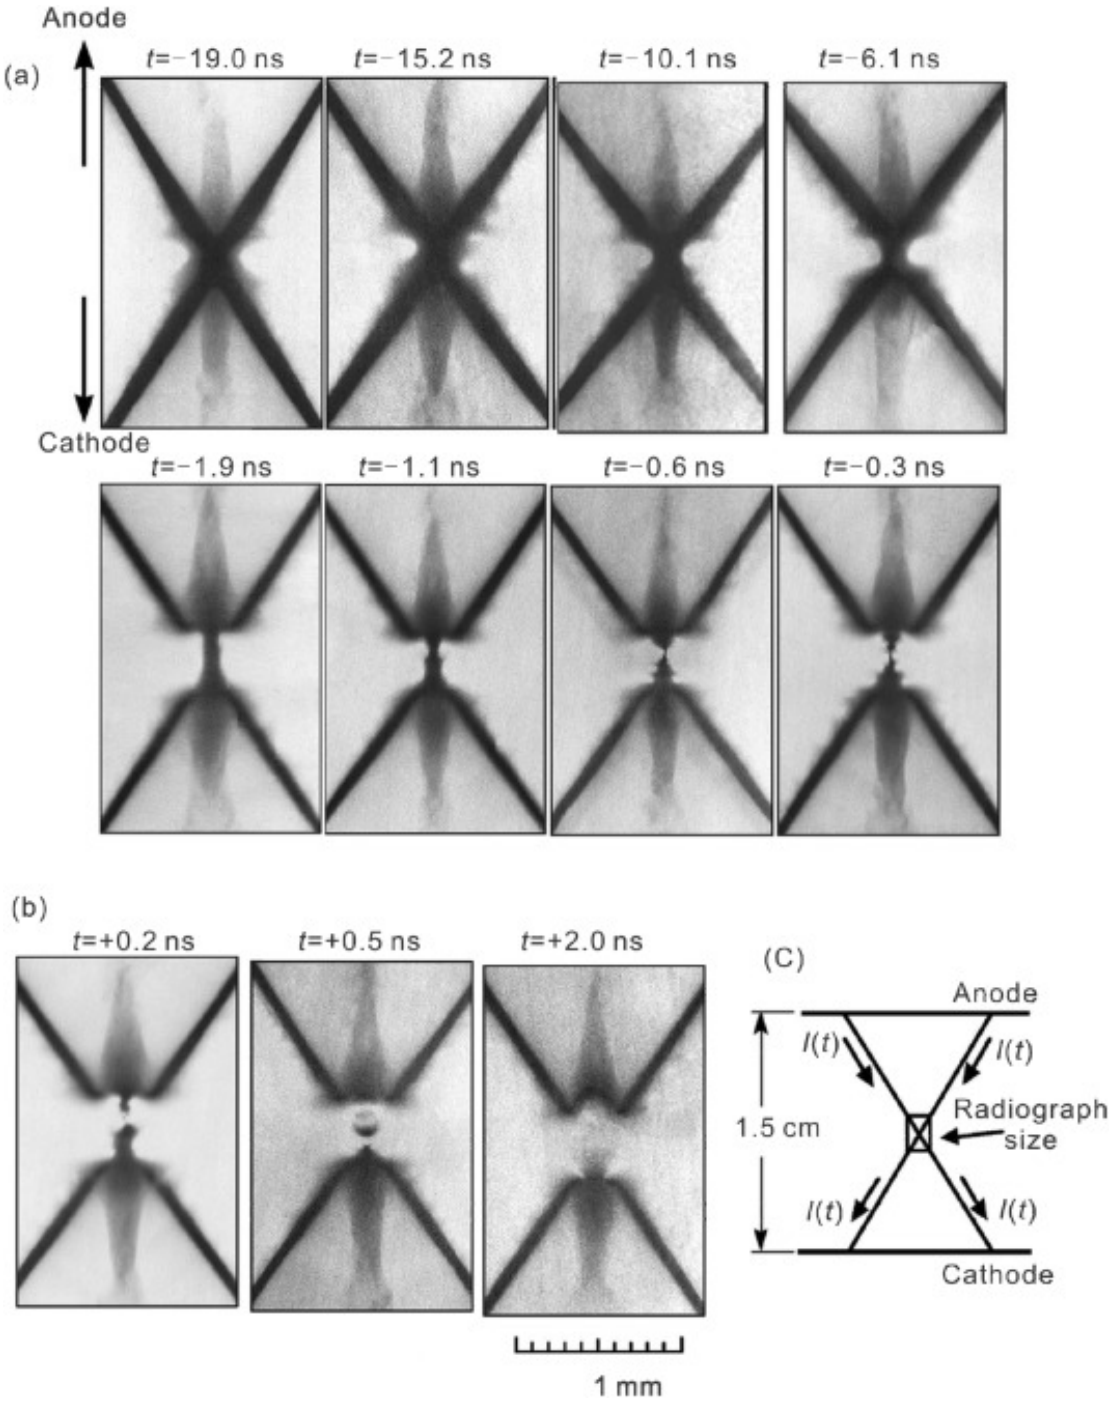


Figure 5.4.: Radiography images of 25 μm X-pinches (top row), and 17 μm X-pinches (middle and bottom rows). Images show a subset of the full system, as demonstrated in the diagram on the bottom right. Time away from target pinch is shown above each image. Reproduced from Ref [21].

Results from radiographs

The global dynamics of 2-wire X-pinches described in §5.1 are informed by radiographs such as these, and therefore do not need to described again here. Instead, key features are pointed out for later comparisons against simulations.

Description: Fig 5.4 shows radiographs of 2-wire molybdenum X-pinches with varying diameter, which were imaged using 17 μm source X-pinches. The dataset for this load material is the most extensive, and is therefore the focus of this section. The spatial resolution for these radiographs are 2-3 μm, limited by the grain size of the imaging plate. The top row of Fig 5.4 show images of 25 μm target X-pinches, as evidenced by the thicker legs compared to other images. The second and third rows are of 17 μm X-pinches.

The values above each image indicates the radiograph time relative to the target's pinch time, obtained by comparing peaks in the PCD signals that correspond to the source/target pinch events. The more similar the mass of the source and target X-pinches, the more closely synchronised their pinch events are. Without the absolute timings, however, it is difficult to determine the current through the load from B-dot measurement. Additionally, the change in inductance from the source X-pinch pinching breaks the assumption of equally distributed current through source and target X-pinches, adding further uncertainty. It is therefore difficult to estimate how much energy has been deposited in the target X-pinch at the time of imaging.

Overall observations: The features are remarkably similar across radiographs given that each corresponds to a different shot, suggesting a highly repeatable system despite the uncertainty in contact point morphology. The system is also symmetric in the left/right direction, although it is difficult to know whether this extends to front/back symmetry, as any variation is integrated in the line-of-sight direction.

Legs: Both the top row (25 μm) and bottom rows (17 μm) show X-pinches that have expanded leg widths compared to their initial radii by 3-4 times, while remaining optically thick to backlighter emission. Neglecting mass loss through ablation, suggests a decrease in density of $\rho \sim \rho_0/10$. As a point of comparison, radiographs of a 12.7 μm Al X-pinch indicate wire expansions of 15-20 times, whereas 13 μm Au wires show an expansion of only 3 times. This variation is expected to be a combined effect of conductivity and heat capacity properties of the material. While the legs are predominantly opaque, they are surrounded by regions of lower density plasma, most visible at the inner ends of the legs where they meet to form relatively flat surfaces, i.e. the 'electrodes' of the mini-diode described previously. Less discernible regions also form along the wire surfaces, for example at -6.1 ns along the lower right leg.

Axial jet: The axial jet is already formed by the time of the earliest radiograph and takes a few more nanoseconds to reach peak (areal) density, which in some images is comparable to that of the legs. In all images it is highly collimated close to the neck. In some cases the jet tapers to a narrow point, in others it expands and becomes unstable at $\sim$0.5 mm away from the neck.

Neck: Of primary interest is the dynamics of the neck itself, which is where strong radiative cooling is expected to take place. By -10 ns, the 25 μm X-pinch has begun to form a micro Z-pinch between 150 to 300 μm long. This column rapidly implodes at 3×10^4 m/s and becomes strongly $m = 0$ unstable. At -0.3 ns it has bifurcated into several smaller regions as small as $7 - 10$ μm, potentially lower. At 200 ps after the pinch, mass is seen to rapidly evacuate the neck region. A shock wave appears to emanate from the neck as evidenced by the hemispherical structures formed within a plasma island at 0.5 ns, and along the top 'electrode' at 2.0 ns. By this point, the plasma density in the neck becomes too low to detect, and no further dynamics within this region are observed.

Further comments

While the experimental data is extensive and gives considerable insight to this complex system, the actual conditions at the time of peak emission still remain unknown. The neck visibly reaches tight compressions, at radii comparable to the resolution of the radiographs, $2 - 3$ μm. High resolution is therefore required to correctly simulate these dynamics, preferably in 2D where high resolution can be used at lower computational expense. However, there are challenges when modelling this system in 2D, which are discussed in the next section.

5.2.2. 2D simulations

While it is desirable to run simulations in 2D to reduce the computational expense, this is particularly challenging for the wire X-pinch due to its inherent 3D structure. Consider an X-pinch with N_w wires and a half-opening angle of θ_w, as shown in Fig 5.5(a). In reality, the WXP has a constant wire radius r_w and density ρ_w, and therefore constant linear mass m_l, along its length:

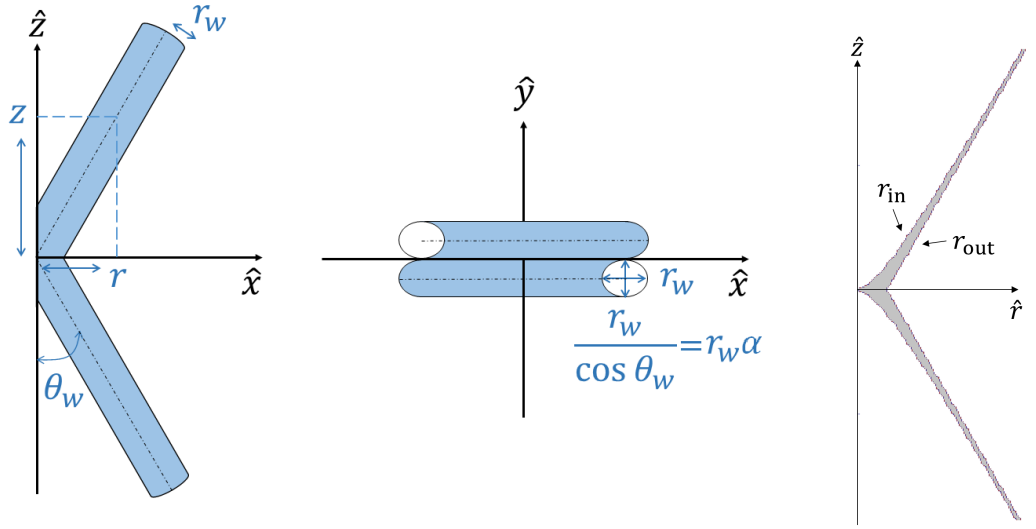
$$m_l = N_w \times \rho_w \times \pi r_w^2 \alpha, \tag{5.1}$$

where $\alpha = 1/\cos(\theta_w)$ is a term accounting for the projection of the wire cross-section on the $\hat{x} - \hat{y}$ plane (Fig 5.5(a), right diagram).

However, in 2D axisymmetric geometry, a constant radius and density corresponds to an annulus with varying linear mass:

$$m_l'(z) = \rho \times \pi \left[(r + r_w)^2 - (r - r_w)^2\right] \times \alpha = \rho \times 4\pi r r_w \times \alpha \ , \tag{5.2}$$

where $r = r(z)$ is the X-pinch radius at a given height z. It is preferable to initialise the simulation by maintaining a constant linear mass along the wire lengths, as the ablation dynamics will otherwise be incorrect. This is only possible in 2D axisymmetric geometry by relaxing the assumption either of constant density or constant radius, as described below.



(a) (b)

Figure 5.5.: (a) Diagrams showing dimensions of an X-pinch in 3D Cartesian geometry, in terms of wire radius r_w and half-opening angle θ_w. Elliptical cross-section in the $\hat{x} - \hat{y}$ plane shown in right image. (b) Profiles of inner and outer surface for 2D axisymmetric geometry when varying radius to maintain a constant linear mass.

1. **Density scaling:**
   Assuming a constant radius, the density can be scaled along the wire by equating $m_l' = m_l$ (Eqs 5.1 and 5.2) such that

$$\rho'(r) = \rho_w \times \max\left[\frac{N_w r_w}{4r}, 1\right], \tag{5.3}$$

   where the multiplicative factor is capped at a maximum 1 so as not to exceed solid density. However, varying the density along the wire also means varying material properties such as pressure and electric conductivities, as shown in Fig 5.7. At initialisation, the variable density at constant temperature results in a high conductivity at the X-pinch crossing point compared to the legs. This causes the current to switch preferentially into the centre, resulting in excessive Ohmic heating and an explosive expansion in the axial direction, which disrupts the hotspot formation.

2. **Radius scaling:**

   The scaling of radius instead of density can be performed in several ways if the profiles of the inner and outer surface can have arbitrary profiles, $r_{in}(z)$ and $r_{out}(z)$, as illustrated in Fig 5.5(b). However, the characteristic global magnetic field is due to the linearly decreasing outer surface, therefore this profile must be enforced, $r_{out}(z) \propto z$. The profiles of the inner and outer surfaces can by derived by first defining the linear mass for arbitrary profiles:

$$m_l' = \rho_w \times \pi\left[r_{out}(z)^2 - r_{in}(z)^2\right]. \tag{5.4}$$

   At the crossing point, the inner radius is chosen to be $r_{in}(r=0) = 0$. Equating Eqs 5.1 and 5.4 therefore gives an outer radius of $r_{out}(r=0) = r_w\sqrt{N_w\alpha}$. Enforcing a linear scaling, the outer radius profile is therefore given by:

$$r_{out}(z) = z\tan\theta_w + r_w\sqrt{N_w\alpha}, \tag{5.5}$$

   and substituting back into the Eq 5.1/5.4 equality gives the following inner radius profile:

$$r_{in}(z) = \left[z\tan\theta_w\left(z\tan\theta_w + 2r_w\sqrt{N_w\alpha}\right)\right]^{1/2}. \tag{5.6}$$

   While the combination of Eqs 5.5 and 5.6 correctly maintains a linear outer profile and enforces a constant linear mass, the difference between the inner and outer radii invariably decreases (with increasing r) to below the cell size. Applying higher resolution with mesh refinement at these regions further away from the crossing point is then counterproductive to resolving the hotspot dynamics.

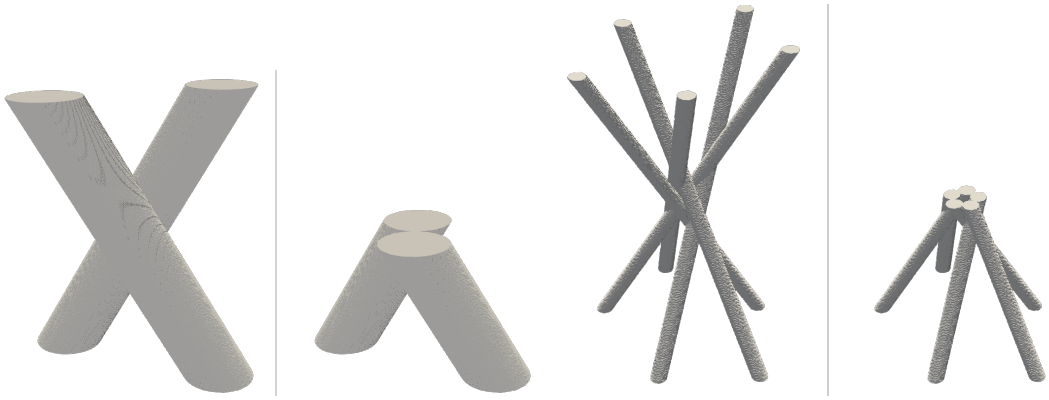
3. **Cartesian simulation:**
   Instead of using axisymmetric geometry, a 2D Cartesian simulation could be run with an in-plane extent (Δy) calculated to match the wire mass:

$$N_w \times \rho_w \times \pi r_w^2 \alpha = \rho_w r_w \Delta y, \tag{5.7}$$

$$\Delta y = N_w \pi r_w \alpha. \tag{5.8}$$

   This would not require any scaling along the wire itself. However, this loses the cylindrical symmetry of the magnetic field and inherent radial convergence, which is important to accurately capture the formation of the hotspot.

While the above represent feasible solutions maintaining a constant linear mass in 2D geometry, the 3D nature of the crossing point cannot be modelled. Consider the cross-sectional views of 2-wire and 5-wire X-pinches, shown in Figs 5.6(a) and (b) respectively. Both are initialised such that their cross-sections in the $\hat{z}$-direction are a ring of circles in contact, with an applied azimuthal skew on each wire to prevent wire crossings. While this initialisation is straightforward to implement analytically, it becomes less physical for a larger number of wires as they will likely be arranged in a non-azimuthally symmetric packing. In the low-wire number limit such as shown in the 2-wire case, the cross-section is irregular and will therefore experience different magnetic field magnitudes parallel and perpendicular to the X-pinch plane. The above gives clear indication that 2D modelling of a fundamentally 3D system will have strong inconsistencies, therefore only 3D simulations of the experiments described in §5.2.2 were performed. These are described in the next section.



(a) 2-wire X-pinch (b) 5-wire X-pinch

Figure 5.6.: Wire X-pinches showing the full wires (left) and the cross-section in the $\hat{x} - \hat{y}$ plane (right), for (a) 2-wire and (b) 5-wire X-pinches.

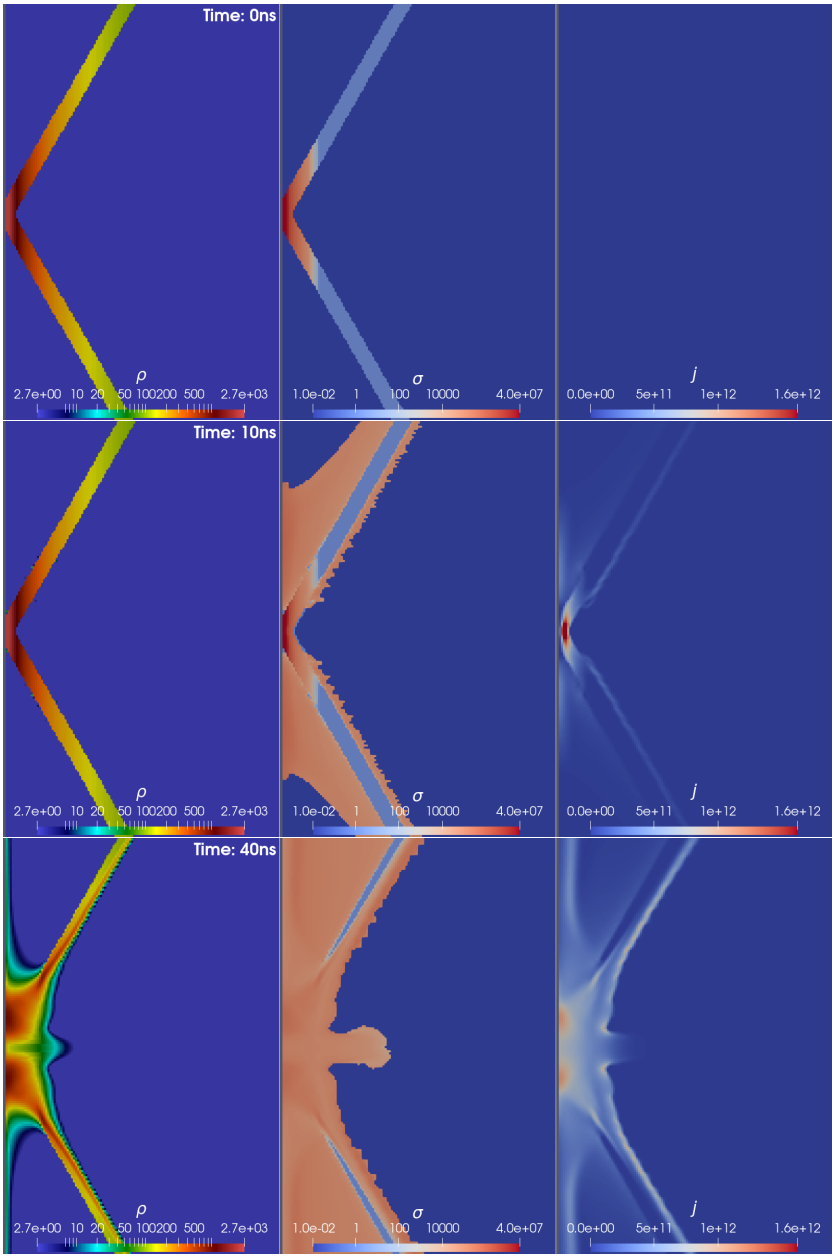


Figure 5.7.: Effect of simulating an X-pinch in 2D axisymmetric geometry, with fixed radius but variable density along the wire to maintain constant mass, at three different times ($t = 0, 10$, and 40 ns). Shown from left to right are profiles of density, electrical conductivity, and current density.

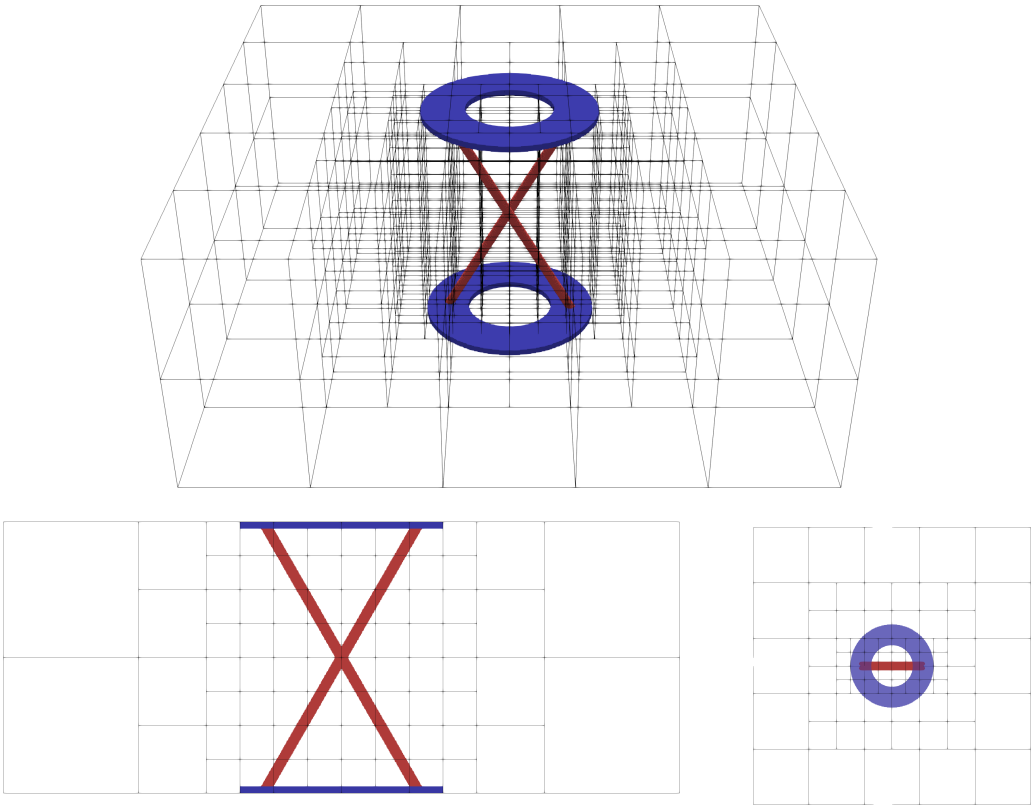


Figure 5.8.: Mesh splitting of a 3D X-pinch simulation, run using a total of 400 processors. Processors from grid splitting shown in black boxes, each of which corresponds to $N_x \times N_y \times N_z = 50 \times 50 \times 50$ cells. The equivalent standard uniform grid simulation at this highest resolution would require 3200 processors. Bottom left and right plots show the grid split in the $\hat{x} - \hat{z}$ and $\hat{x} - \hat{z}$ planes respectively.

5.2.3. 3D Simulations

Previous simulation studies of X-pinches in 3D showed low density coronal plasma being rapidly ejected and hitting simulation boundaries, leading to numerical instabilities [113]. Increasing the domain size to account for this required running at too low a resolution (10 μm) to sufficiently capture hotspot dynamics. As demonstrated in the single wire case in §4.5, the static refinement capability allows for substantially larger boundaries at a minimal computational expense, and high resolution to appropriately capture the small-scale dynamics.

Set-up

The domain extents used in this simulation was $L_x \times L_y \times L_z = 2 \times 2 \times 0.8$ mm^3. The base grid was resolved at 16 μm, while the region around the pinch was refined at 2 μm. The mesh splitting used is shown in Fig 5.8. Note, the higher resolution region has different extents in the $\hat{y} - \hat{z}$ and $\hat{x} - \hat{z}$ planes as seen from the top-down view, in order to minimally but completely resolve the X-pinch wires.

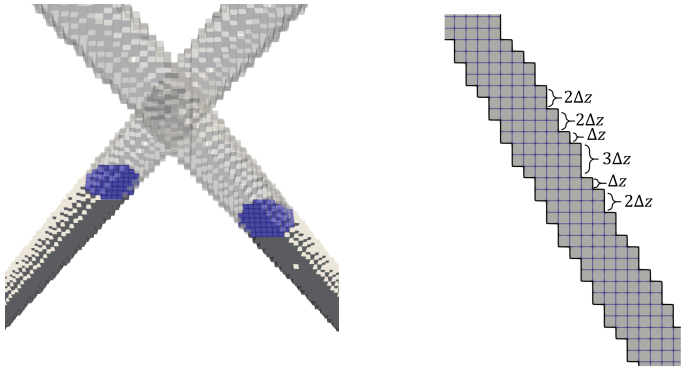


Figure 5.9.: (a) Initialisation of 17 μm wires on a 2 μm resolution grid. Cross-section of discretised wires uses $\sim 8 \times 10$ cells. (b) Staircasing effects due to discretising angled features on Cartesian grids, leading to irregular surfaces.

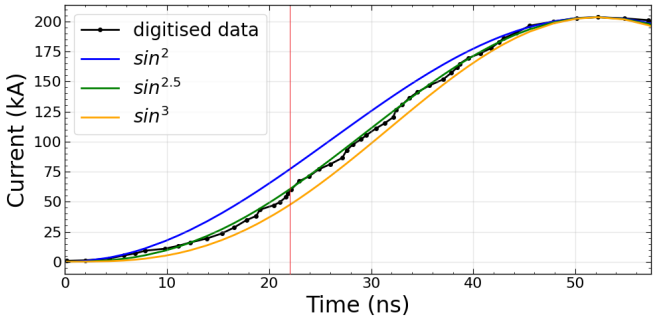


Figure 5.10.: Matching the experimental load current from B-dot probes to sinusoidal profiles of different powers. Vertical red line indicates $t = 22$ ns.

For comparison against experiments, the simulation below is of a 2-wire, 17 μm molybdenum X-pinch. Using the studies performed in §4.3 to inform choices, the models used were FEOS for equation of state, LMD for transport, and the group-resolved loss for radiation. As described in §5.3.2 for tungsten, copper conductivity tables and atomic data were used with relevant scaling factors applied.

The wires were initialised on the grid by building up the cross-section in the axial direction. For vertical wires (as in §4) the cross-section is circular and for angled wires the cross-section is elliptical, as shown in Fig 5.9(a). For a 17 μm wire initialised on a 2 μm resolution grid this corresponds to roughly 8×10 cells. Another consequence of discretising angled wires is the 'staircasing' effect, shown in Fig 5.9(b). The size of the axial 'steps' as a function of radial position are not constant, which consequently leads to unphysical variation in the current density profile along the wire surface.

Also described in §5.3.2 is the sensitivity of simulation dynamics to the current profile, which was retrospectively realised to make a substantial difference to implosion timings. As discussed below, the system reaches a compressed state much earlier than radiographs in Fig 5.4 seem to indicate. This simulation was run using a commonly used $\sin^2$ approximation to the current pulse, however the digitised current pulse shows better agreement with a $\sin^{2.5}$ profile. While this may appear as a small difference, the

difference in total magnetic energy deposited at 22 ns ($E \sim \int_{t=0}^{t=22} I^2(t)$) is considerable; $E(I \sim \sin^2)/E(I \sim \sin^{2.5}) = 1.83$, i.e. almost double.

Simulation dynamics

The evolution of a typical 3D simulation is shown in Fig 5.11, until 22 ns when the simulation was terminated. These figures were produced by taking a slice through the $\hat{y} - \hat{z}$ plane at $x = 0$, shown relative to the 3D 'electrode' rings. As such, the details of 3D effects are not captured and will be presented next. Note, these electrode rings are used to prevent discontinuities in magnetic field and therefore currents forming on the boundaries, while having a hollow centre allowing mass to be transported out of the system from jetting. For reference they are initialised with a 0.3 mm outer radius in this simulation.

5 ns: At early time the core-corona structure is beginning to form, and the central slice shows a slightly denser region through the front wire and a less dense contribution from the rear wire. Axial motion from the low density material at the crossing point is apparent although not fully developed.

7 ns: Two nanoseconds later, the X-pinch legs have continued their expansion, while the central crossing point remains at high density. This is due to the inner surface of each wire being shielded from the magnetic field by the opposite wire, undergoing reduced Ohmic heating and ablation. The overlap region of the two angled wires as viewed from the $\hat{y} - \hat{z}$ plane forms a rhombus, which is reflected in the profile of the high-density, shielded region. There is also an obvious axial jet feature at densities between $10^{-3} - 10^{-2}$ kg/m^3 which is rapidly ejected at 2.7×10^5 m/s, similar to the early time axial jets observed in radiographs. At this stage, the jet is formed in near equal contribution due to the normal component of the wire ablation colliding on axis, and the pressure gradient at the crossing point which is untamped due to magnetic nulls.

Additionally, ablation streams along the wire lengths exhibit modulations with wavelengths of the order of $10 - 15$ μm. The wavelength of these modulations suggests it is not the staircasing effect associated with discretising wires on a Cartesian grid (which has an expected variation between $1 - 3\Delta z = 2 - 6$ μm), instead this perturbation appears to originate from the 'fundamental mode' of wire ablation as described in the previous chapter. Finally, although the ends of the wires remain in contact with the electrode, there is enhanced ablation around this region which expands radially outward.

10 ns: By this time, the wire cores (defined as $\rho = \rho_0/1000$, see discussion of Fig 4.7) have expanded further, from 52 μm at 7 ns to 106 μm at 10 ns. This rate of expansion 1.8 cm/μs is a factor of 2.5 greater than the 0.72 cm/μs observed in Al single fibres in §4. This can be explained by the difference in electrical conductivities at solid densities; aluminium is almost twice that of molybdenum. The resistive diffusion timescale $\tau \sim r_0^2/\sigma$ is therefore longer for molybdenum, allowing more energy to be deposited in the wire. The characteristic rhombus shape of the high density regions

intersect persists, and grows along with wire cores. Low density material has entirely filled the regions above and below the wires, but not to the sides. The $\mathbf{j} \times \mathbf{B}$ force, enhanced by the global magnetic field contribution around the wires acts to strongly tamp the radial expansion.

15 ns: The core continues its near-constant expansion rate for a further 5 ns, by which point it has expanded to 180μm. This represents a 10$\times$ expansion in radius. The states along the legs appear smooth, however the modulations observed in the ablation streams at 7 ns imprint on the core, giving rise to banded modulations formed along the wire core. The jet density increases, and forms a hollowed structure. This is due to the jet carrying a significant current, which Ohmically heats the increasingly dense material, raising its pressure and forcing an expansion outwards that resists the ram pressure from the inflows. This is in contrast to the jet behaviour in radiographs, where the jet appears narrow and highly collimated.

20 ns: By this time the wires, central core, and axial jets are disrupted. The $\mathbf{j} \times \mathbf{B}$ force that acts to confine ablated wire material on the outer surface begins to exceed the thermal pressure gradient and drive a compression through the material normal to the surface, causing a density pile-up. The enhanced radial Lorentz force at the neck begins to compress it strongly, initiating the formation of the micro Z-pinch. There is a skew on the slight radial jetting due to 3D effects, discussed below. The axial jet disintegrates as it exits the simulation domain. This behaviour bears resemblance to experimental radiographs, where the axial jets appear collimated up to a certain extent then abruptly turn turbulent. Also close to the electrode rings is evidence of enhanced current on the wire core. This heats the region in contact with the electrode, while launching another compression axially inwards.

22 ns: Two nanoseconds later the core reaches a highly compressed state and forms a clear Z-pinch structure. The radially inward velocity between 20 ns to 22 ns is 50 μm/ns, faster than the 30 μm/ns observed in experiments. The states in this region are discussed below. The second difference is the apparent collimation of the jet, which forms a thin, unstable structure within a wider also coherent one. The recovery from the disintegrated state is due to continued ablation off the wire surface increasing the jet density, causing further cooling.

Finally, the simulation is terminated due to low density plasma hitting the magnetic field boundary at the x and y domain edges. The forced magnetic field boundary value is inconsistent with that carried by the plasma, forcing unphysical currents. The two ways in which this could be overcome is either by restarting the simulation into a larger extent region and letting the simulation continue further, or by restarting from $t = 0$ with a larger domain. The former is feasible but not straightforward, as the processor naming convention uses a space-filling curve (see §3.1) and a restart capability is more complicated than identifying processor regions from file names. The latter is straightforward yet time-consuming, and will be performed in future simulations to investigate the hot-spot even later in time.

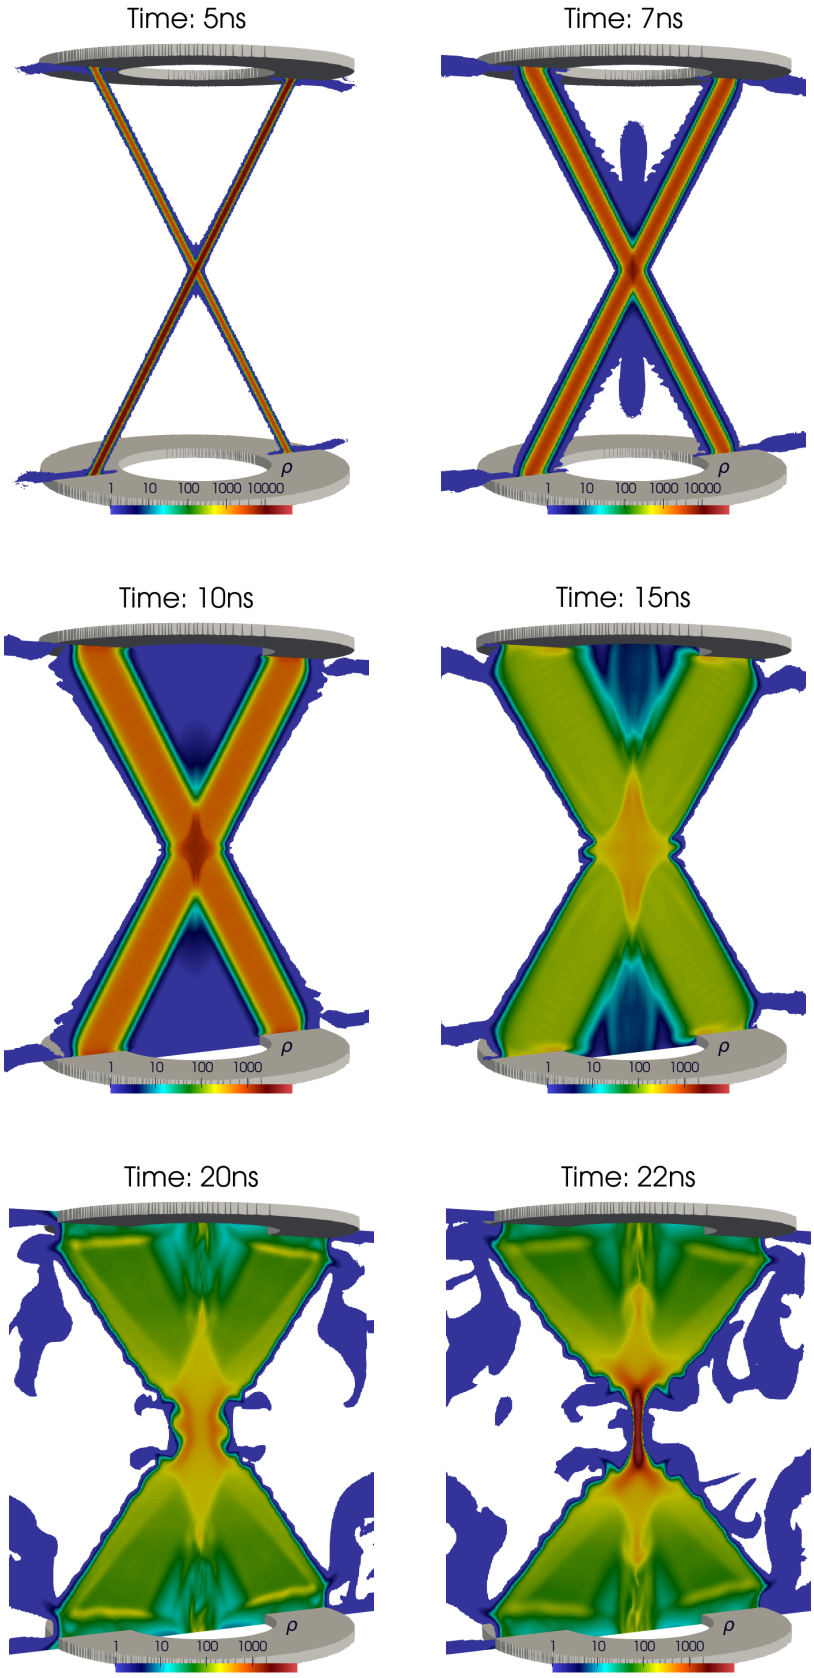
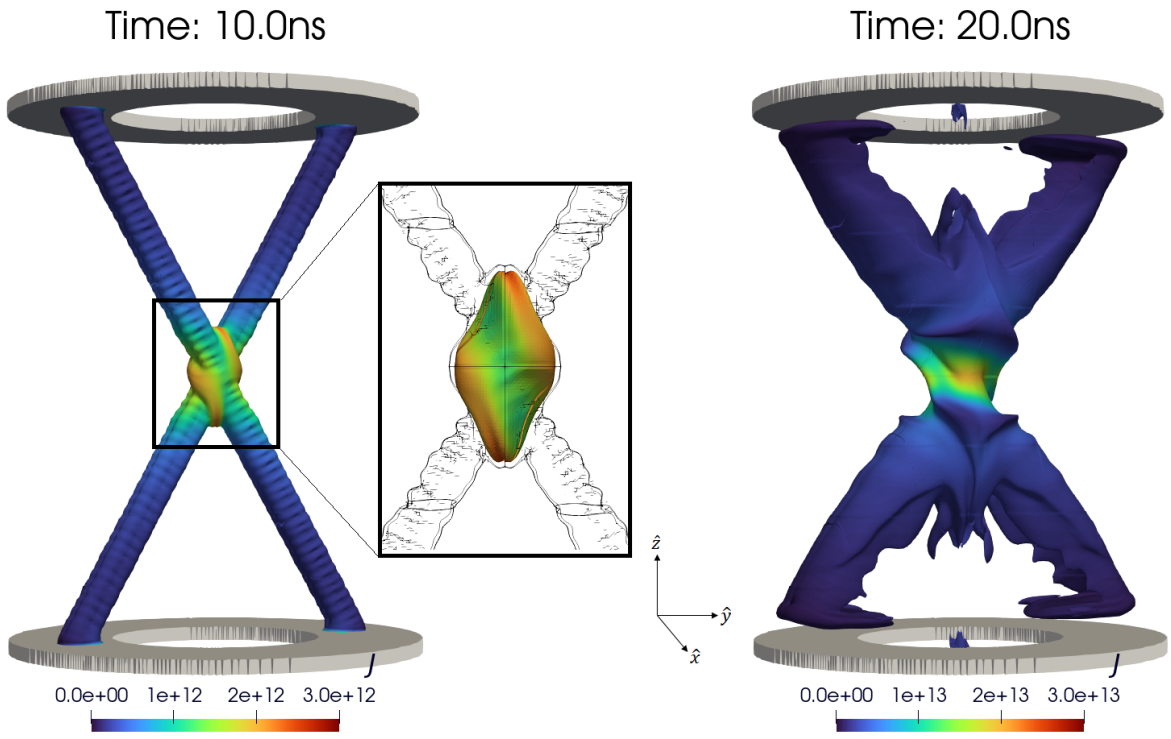


Figure 5.11.: Evolution of 3D simulations of a 2-wire X-pinch, shown by taking slices through the $\hat{y}-\hat{z}$ plane of density output, at six different times. Electrode rings shown in grey for clarity

3D structure

While the central slice appears symmetric in the left-right direction, there is significant 3D asymmetry in regions away from the central plane, due to the wire centres being offset from the axis by their radii (visible in Fig 5.6(a)). This asymmetry is demonstrated in Fig 5.12 at 10 ns and 20 ns using density iso-contours. At early time, the $\rho = 500$ kg/m^3 surface along the wires is corrugated with the same wavelength as the low density ablation seen in Fig 5.11 at 7 ns. Material from the crossing point expands in regions where it is not shielded from the magnetic field by the wires, forming a twisted rhombohedron shape (inset in Fig 5.12(a)). This corresponds to the high density rhombus observed in the 2D slices above. The exposed regions encounter a high current density as they continue to expand, maintaining a similar shape through the dynamics. At 20 ns, the expansion is apparent in the $\hat{y} - \hat{z}$ plane (along the plane of the X-pinch), while the perpendicular plane ($\hat{x} - \hat{z}$) remains compressed, indicating different rates of implosion in the two directions. However, without experimental radiography data from a side-on view it is difficult to determine if implosion asymmetry is physical. By 22 ns, the micro Z-pinch is a nearly uniform column with equal radii in the $\hat{x}$ and $\hat{y}$-directions. The expanded 'electrode' material where the legs meet (see discussion of experimental radiographs in §5.2.1) extends further outwards in $\hat{x}$-direction (not visible from this



(a) $\rho = 500$ kg/m^3 (b) $\rho = 100$ kg/m^3

Figure 5.12.: Density iso-contour surfaces showing 3D effects in 2-wire X-pinches, shown at (a) 10 ns for 500 kg/m^3 and (b) 20 ns for 100 kg/m^3. Inset in (a) shows twisted rhombohedron surface where crossing-point material is not shielded by the wires, allowing it to undergo Ohmic heating and expansion.

end-on view) than in the $\hat{y}$-direction visible in Fig 5.12. These 2-wire systems clearly exhibit complicated 3D dynamics due to the wire radius offset that cannot be replicated in 2D simulations or by assuming symmetry in the left/right and top/down directions as in previous simulation studies such as Ref [114].

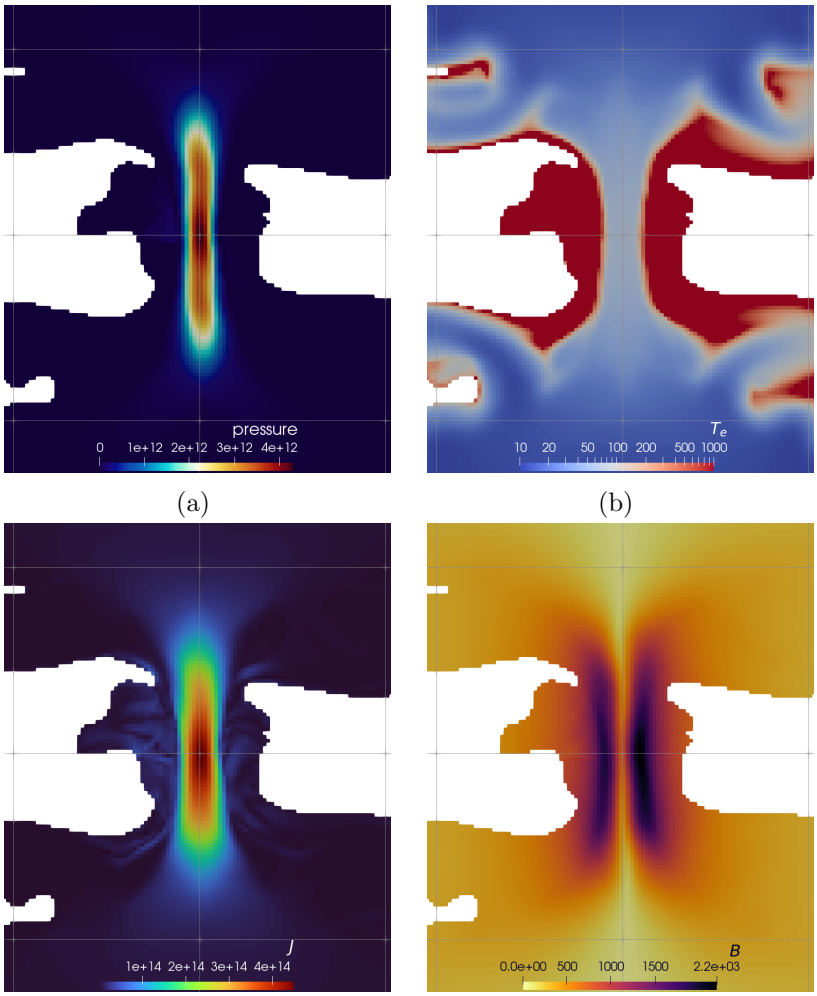
Hotspot

By 22 ns the neck has formed a single high pressure hot-spot, as shown in Fig 5.13 in terms of pressure, electron temperature, current density, and magnetic field. The first and third figures show an elongated region approximately 20 μm wide and 160 μm long. As mentioned above, the width is also 20 μm in the perpendicular plane, not visible in these images. The density in this region is comparable to solid at $\geq 10^3$ kg/m^3. At the centre is a visible peak in pressure which reaches 4TPa and 7.3×10^3 kg/m^3, equivalent to $n_e = 3.6 \times 10^{23}$ cm^3, an electron density similar to spectroscopic analysis of previous similar systems (see §5.1.2). The electron temperatures reached at the centre of the hotspot are relatively cool (<100 eV), however rapidly increase to several keV in the low density coronal plasma. Up to 1 ns before the hotspot formation the current was carried mostly by the corona, by 22 ns it has switched entirely into the core.

The micro Z-pinch shown in these images bears resemblance to the single wire simulated in 3D and presented in the previous chapter, particularly when the current switches into the wire core (30 ns). The $m = 0$ instability forces the wire to pinch and form hotspots along its length, which similar to the neck of the X-pinch, exhibit high pressure and density regions in the centre at relatively low temperatures compared to the corona. The smallest hotspot observed in the single wire case reached 10 μm radius, half the size of the X-pinch hotspot at 22 ns, although simulated at half the resolution. After 30 ns, the wire hotspot was shown to rebound and undergo a second phase of hotspot formation at 35ns as a result of the $m = 1$ instability, which reached pressures up to $3 - 4\times$ the pressure from the first phase.

While the 4 TPa reached at the centre of the X-pinch is considerable, and over an order of magnitude greater than the pressure in the wire ($m = 0$) hotspots, the above suggests that progressing this simulation further could allow the neck to undergo $m = 1$ instabilities, resulting in even further compression. This would likely require a simulation run at 1 μm resolution or smaller to accurately model the compression across multiple cells, and a larger simulation domain to account for expanding material hitting the domain boundaries. Both are feasible using mesh refinement, but remain challenging mostly due to timestep considerations.

The next section describes simulations of a different type of X-pinch, called the solid X-pinch. This configuration uses the same focussing effect to reach extreme conditions in the central point, but does not exhibit the same asymmetry in 3D and can therefore be run at high resolution in 2D.



(a) (b)

(c) (d)

Figure 5.13.: Hotspot conditions in a 2-wire X-pinch at 22 ns, showing 2D slices in the $\hat{y} - \hat{z}$ plane of profiles of (a) pressure in Pa, (b) electron temperature in eV, (c) current density in A/m^2, and magnetic field in T.

5.3. 5 MA solid X-pinches

This section investigates experiments from Ref [22], performed on a 6 MA generator which corresponds to a 30× increase in peak current delivered to the load, compared to the current per X-pinch from the previous section. These experiments are the highest current X-pinch studies conducted at the time of writing, and are the subject of the remainder of this chapter.

5.3.1. Review of experimental results

Machine and load

The driver in these experiments was the SATURN generator at Sandia National Laboratories, a facility typically used for generating hard x-ray bremsstrahlung radiation. SATURN has a peak current of 6 MA and rise time of 60 ns, corresponding to an average $\dot{I} \sim I_{peak}/\tau_{rise} = 10^{14}$ A/s, comparable to the Z generator where $\dot{I} \sim 20$ MA/20 ns$= 2 \times 10^{14}$ A/s.

The crossing point linear mass required to match this current was calculated using a self-similar implosion scaling law, derived first by D. Ryutov in Ref [115]. This considers the simplified equation of motion of a cylindrically imploding annulus with linear mass m_l:

$$\frac{m_l}{2\pi r}\ddot{r} = -P_{mag} = -\frac{1}{2\mu_0}\left(\frac{\mu I}{2\pi r}\right)^2, \tag{5.9}$$

which can be rearranged and expressed in terms of dimensionless quantities $\tilde{r} = r/r_0$ and $\tilde{I} = I/I_{peak}$ to give

$$\tilde{r}\ddot{\tilde{r}} = -\Pi\tilde{I}^2, \tag{5.10}$$

where Π is a scaling parameter used to maintain the same implosion behaviour for different peak currents and initial radii (r_0):

$$\Pi = \frac{\mu_0 I_{peak}^2 \tau_{rise}^2}{4\pi m_l r_0^2}. \tag{5.11}$$

While the expression above was intended for liner implosions it has also successfully been applied to X-pinch loads. In this case, m_l corresponds to the sum of individual wires and r_0 is the crossing point radius, which for closely packed wires is close to $\sqrt{N_w} r_w$ (where r_w is the wire radius).

The above expression was used to scale X-pinch loads from 1 MA experiments conducted on the COBRA facility at Cornell University [116]. These experiments characterised large wire number (32-64 wires) X-pinch loads with linear masses around $m_l = 3$ mg/cm. Scaling this value to 6 MA gives a required linear mass of 18.1 mg/cm, therefore a range of values around this estimate were tested in the SATURN experiments (12-24 mg/cm). Wire X-pinches were initially tested using up to 128 wires to achieve

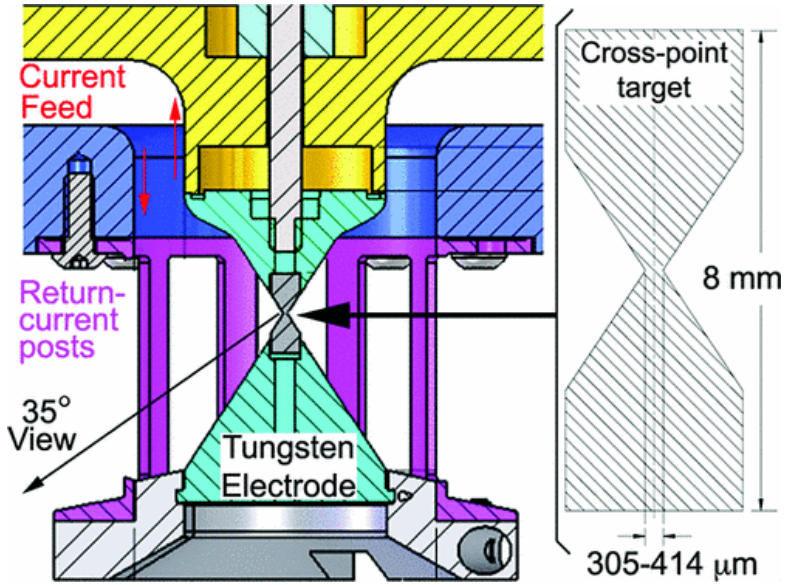


Figure 5.14.: CAD model of load hardware, showing electrodes in cyan and solid X-pinch load in grey etched. Dimensions of load are shown on the right. Note, the 35° diagnostic viewing angle. Reproduced from Ref [22].

this linear mass. However, issues were encountered with reproducibility, wire packing at the crossing point, ease of manufacturing, etc.

An alternative load was developed referred to here as the solid X-pinch (SXP), shown in Fig 5.14. These are an extension of the HXPs described in §5.1. Instead of a thin wire connecting separate conical electrodes, SXPs are manufactured from a single block. This has the advantage of reproducibility and cylindrical symmetry, however is only suitable for high current generators that require relatively large crossing point radii (100s of μm) due to fragility considerations. Several load materials were tested of which pure W and W-Cu alloys produced the highest x-ray yields, however for simplicity only the former is considered in subsequent discussion.

Diagnostics

Several diagnostics were fielded on this experiment, further details can be found in Ref [117]. As with simulation studies in previous sections, only a few specific ones are described here, along with key findings from each that are used for later discussion pertaining to simulation results.

- Time-integrated slit imaging
  This diagnostic passes emission through microfabricated slits of varying width [118] to estimate the source size. For small source sizes or small wavelength emission, the resulting image will look like standard point-projection. For source sizes of order slit width or for longer wavelength emissions, the system is diffraction limited and an interference pattern is produced. Wave-optics simulations can be performed assuming different spot sizes for a given slit width and wavelength,

and compared against experimental data to estimate the source size. In this experiment, slits with 20, 40, and 80 μm widths were fielded with 40 μm Cu filters that transmit in the 6-9 keV range (Fig 21 in Ref [117]). Comparison against wave-optics modelling suggests the smallest source size obtained across all experiments was a 5 ± 2 μm source from a tungsten load (S3941).

- Time-resolved XRDs
  The power and energy from soft X-ray emission ($0.1 - 1$ keV) was measured using x-ray diodes (XRDs), and hard emssion (>1keV) was measured using photoconducting diodes (PCDs). Example XRD traces are shown in Fig 5.15 relative to current measurements from B-dot probes for four experiments with different load materials. All signals show between 2-3 distinct peaks and several smaller ones in the x-ray emission, suggestive of multiple pinch events. The two Cu-alloy experiments only produce signals after 60 ns, whereas the pure tungsten loads have earlier bursts near 50 ns. These times are considerably earlier than peak current, implying non-optimal coupling of energy from the generator into the X-pinch. Ref [22] notes that the diodes had a temporal resolution of 200 ps which might be insufficient to fully capture the burst event.

- Time-gated pinhole imaging
  This diagnostic passed emission through an aperture, which was filtered into monochromatic soft emission and broadband harder emission. The soft 277 eV emission was obtained using a multi-layer mirror (MLM), where only wavelengths that satisfy a Bragg angle are reflected into the detector. The higher energy emission (>1 keV) was obtained using Be and CH filters. Fig 4 in Ref [22] shows time-gated emission at these two frequency ranges as a function of time. Close to peak burst (as timed against XRD traces), the soft emission appears as a large spot roughly 2×3 mm^2, and the harder emission appears as a narrower band within the soft emission region. This is in contrast to typical WXP behaviour, where the soft emission occurs from small localised hotspots and the harder emission occurs from accelerating electrons across a larger gap. Further discussion of this diagnostic is presented in §5.3.4 along side synthetic self-emission images.

- Time-resolved x-ray streak camera
  This diagnostic uses slit imaging with various filters along with a streak camera to get time resolved spectral intensity measurements. An example of the raw data is shown in Fig 5.16. Ignoring misalignment effects, the time axis is on the bottom while the spatial axis on the left shows variation of intensity using different filters. The streak camera spectral response to the different filters is also shown in the adjacent image. The relative intensities from a lineout of the streak image is shown in the bar chart at the bottom of Fig 5.16, shown in red lines. A corresponding predicted measurement can be made by assuming a spectral profile and multiplying it with the filter spectral response. This was performed for

multiple continuum spectra, and agreement with measurements were found using profiles from two sources; a 1275 keV blackbody, and a 3500 keV bremsstrahlung source. The close match of each with experiment makes it difficult to determine which is more representative of the experimental conditions. This is also discussed further in §5.3.4.

Overall, the data collected points towards extreme conditions that were reached through unique means. As a point of comparison, the best performing wire array experiments on SATURN produced similar values of peak x-ray power per unit length (4 TW/mm for wire array, 3 TW/mm for solid X-pinch). This was highlighted by the authors as being of considerable interest, since in the former case the energy conversion to x-rays comes from the wire array kinetic energy, while the SXP is stationary. Simulations are a necessary tool for understanding diagnostically challenging experiments as these, and the first detailed numerical study of this system is described next.

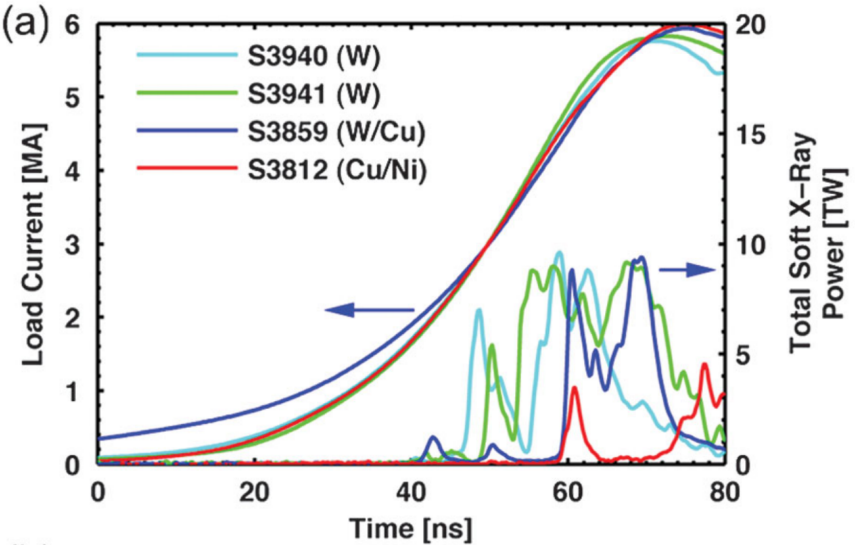


Figure 5.15.: Experimental current traces obtained using B-dot probes for four SATURN SXP experiments using different load materials. Overlaid are soft x-ray (< 1keV) traces obtained using XRDs, corresponding to the right-hand axis. Reproduced from Ref [22].

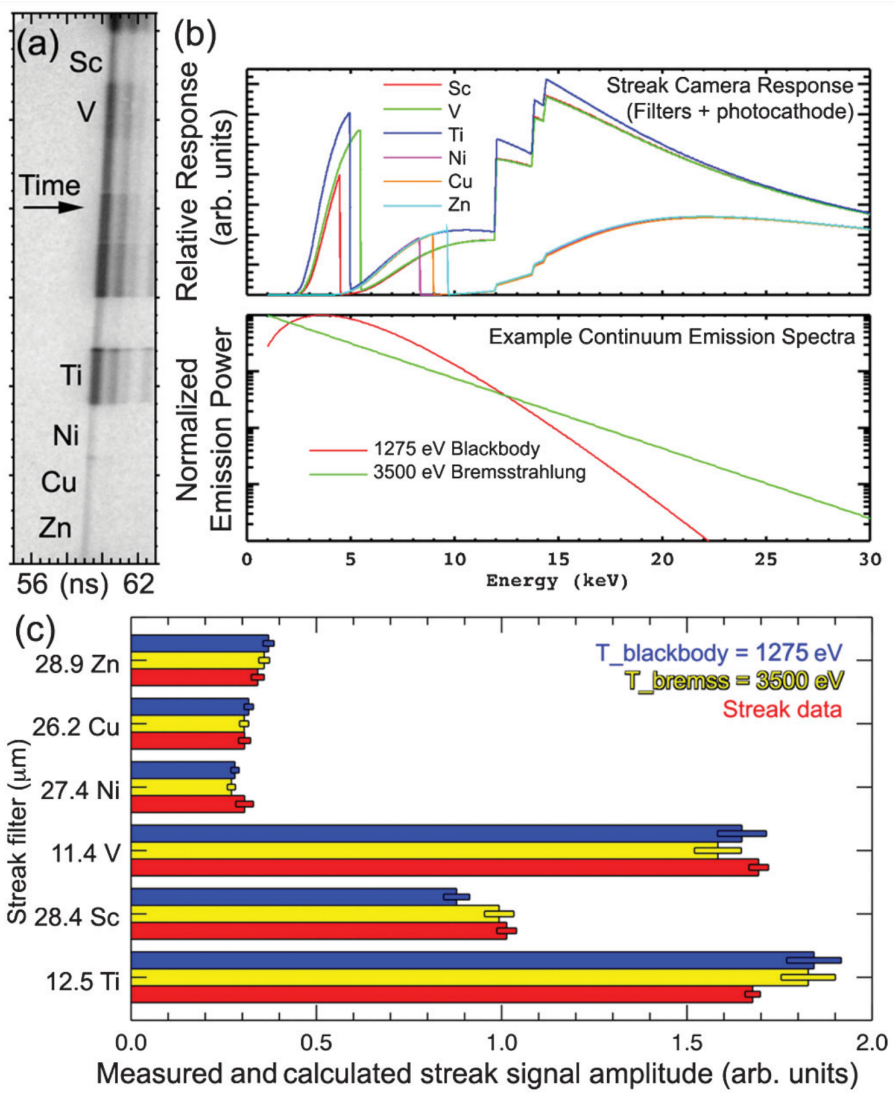


Figure 5.16.: X-ray streak camera imaging of an SXP load. (a) Raw streak images of emission through a slit using ten filters, between 56−62 ns. (b, top): Filter responses for six of the ten filters. (b, bottom:) spectra from a 1275 eV blackbody source and 3500 eV bremsstrahlung source. (c) Experimental intensity amplitudes from the six filters, and corresponding calculated amplitudes using the above spectra. Reproduced from Ref [22].

5.3.2. Simulations

Set-up

Simulations of tungsten SXPs fielded on SATURN were run in 2D axisymmetric $(r-z)$ geometry as the configuration has inherent cylindrical symmetry, thereby sidestepping the issues described in §5.2.2. The domain extents used were $L_r \times L_z = 4 \times 2$ mm^2, and mesh refinement was used to resolve the axis at $\Delta r = \Delta z = 1$ μm and 8 μm at the far edge, as shown in Fig 5.17.

The SXP was initialised with a 398 μm diameter neck, corresponding to the highest value of linear masses tested and tungsten solid density ($r_0 = \sqrt{m_l/\pi\rho_0}$), with an opening angle of $\sim 28°$, shown in dark grey in Fig 5.17) FEOS equation of state and LMD transport tables were used instead of analytic models, see §4.3 for relevant discussion. As a surrogate for modelling tungsten loads, copper LMD tables were used with density and conductivity scaling factors corresponding to the ratio of these parameters at ambient conditions (ρ_0^W/ρ_0^{Cu} and σ_0^W/σ_0^{Cu}). Similarly, copper opacity and emissivity tables used with an applied scaling factor of Z^W/Z^{Cu}. Unlike previous simulations of thin wire systems, a higher vacuum density threshold was of 10^{-2} kg/m^3 was used to avoid the timestep restrictions encountered at high Alfvén speeds. The choice of threshold value is discussed further below within this section.

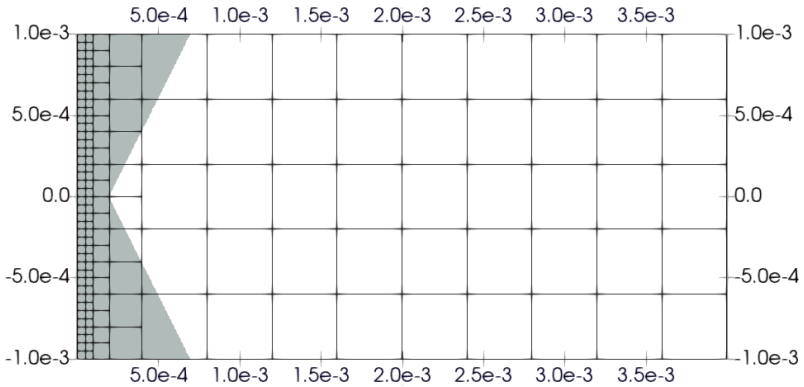


Figure 5.17.: Set up of 2D axisymmetric simulation, showing the SXP in dark grey and with grid splitting in black lines. The splitting shows processor boundaries, each of which has 50×50 cells, therefore smaller boxes correspond to smaller cell sizes.

Effect of current profile

The dynamics of this system were found to be highly sensitive to the current profile used. Simulations were initially run with a $\sin^2$ current profile, a commonly used approximation for generators like MAGPIE and Z. However, simulations were consistently found to reach peak compression at ~ 30 ns, less than half the rise time. It is clear to see why this was the case from Fig 5.18, which compares the experimental current profile obtained from B-dot probes against sinusoids to different powers:

$$y = I'_{peak} \sin^n \left(\frac{2\pi t}{4 \times \tau'_{rise}} \right), \tag{5.12}$$

where I'_{peak} and τ'_{rise} are the peak current and rise time extracted from the digitised profile of S3940 in Fig 5.16. The $\sin^2$ profile has a significantly faster dI/dt at early time, reaching over twice the experimental current by 30 ns. A better fit is seen for $\sin^5$, however this underestimates the current at early time and overestimates it at later times. Instead of fitting an analytic profile the current data was read directly into the code and linearly interpolated between points. This resulted in a pinch time of 57.6 ns, which is much closer to the peaks in XRD signals for the same shot (assumed to be pinch events).

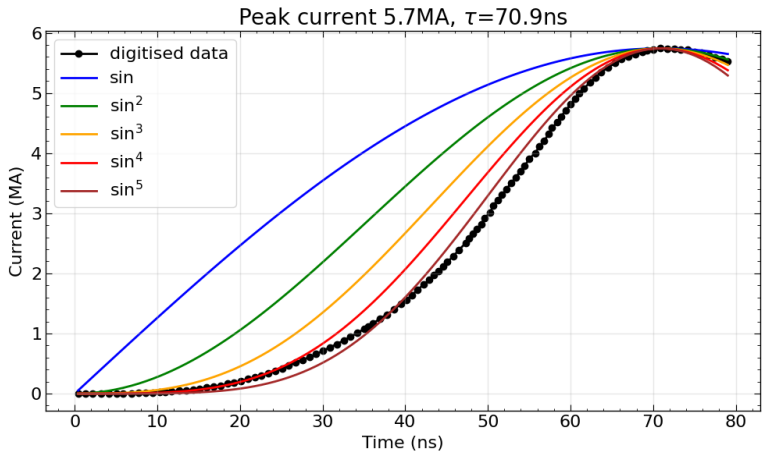


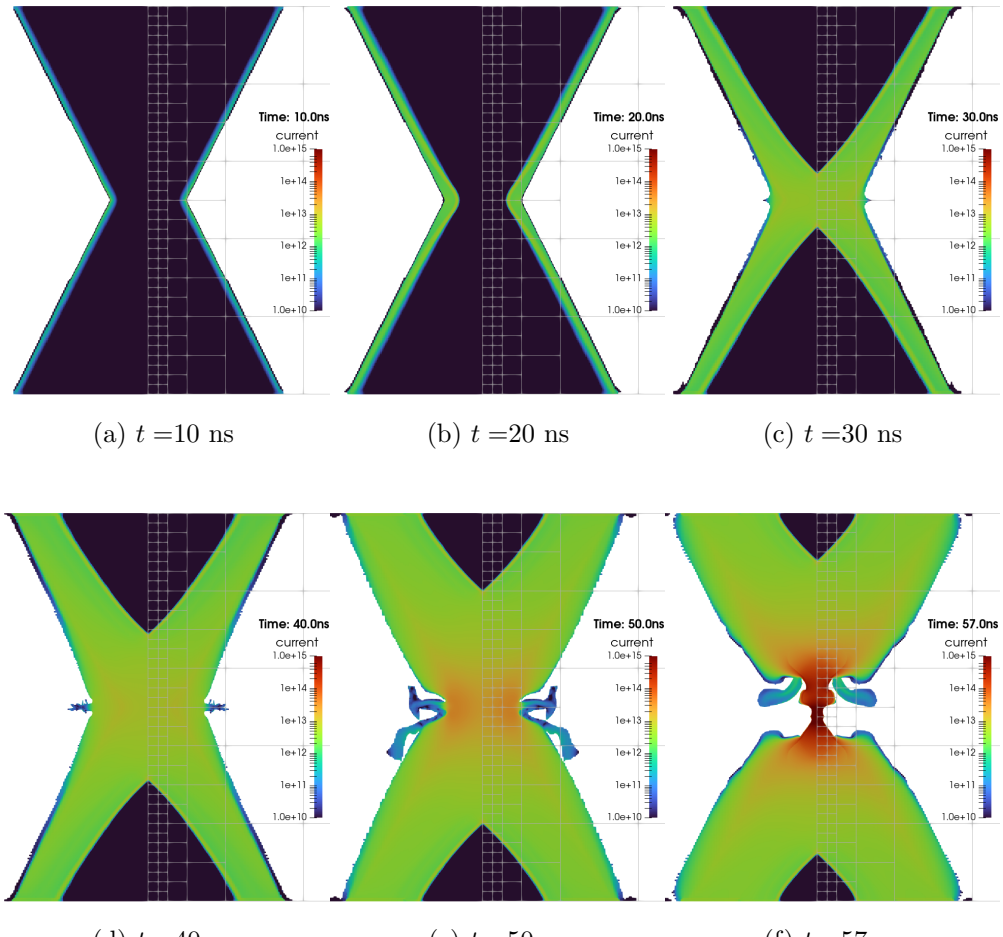
Figure 5.18.: Matching load current data to $\sin^n$ profiles, for varying n. Using $\sin^2$ will give excessively high current at early time compared to data. Other powers also are only a good fit over short intervals. Instead, interpolated fit of data points was used.

Both the experimental and simulation pinch times are 10-15 ns away from peak current, despite using a higher value from the range of linear masses tested. The magnetic energy deposited in the SXP at 57.6 ns as a fraction of the total current pulse is $\Delta E_{mag} = \int_{t=0}^{57.6} I^2(t) / \int_{t=0}^{\tau_{peak}} I^2(t) = 0.32$ using $I(t) \sim \sin^5$, suggesting a significantly under-massed load and non-optimal energy coupling. It is unclear why this is the case, and future work will investigate the effect of scaling load density or crossing point radius on time of peak compression.

Overall dynamics

The evolution of a typical simulation is shown in Fig 5.19, run with the default improved radiation loss model. The early time behaviour is fairly insensitive to the choice of radiation model, while the late time behaviour shows more sensitivity and is discussed further in §5.3.3. Note, for ease of visualisation most results in this section show the output reflected along the z-axis, and the grid splitting overlaid over the actual simulation domain.

Figs 5.19(a)-(c) show the current density slowly diffusing through the solid, conductive material until it reaches the axis at $\sim$ 27 ns. As the current density propagates it Ohmically heats the material, dropping the conductivity ($d\eta/dT > 0$ in the solid phase), allowing it to be driven further inwards. Once the wave reaches the axis it begins to pile up in the centre, as seen at 40 ns. Some radial jetting is observed, however this is at significantly lower density relative to the bulk material and carries much less current. Despite a large current density the SXP does not exhibit bulk hydrodynamic motion until 45 ns, when the crossing point begins to pinch inwards and forms a visible neck. By 57 ns, the neck begins to bifurcate at two or three points. While the low density radial jetting does not affect bulk motion, it creates small perturbations that dictate where along the neck the hotspots form. In this case, the jetting is swept upwards towards the top half of the neck which leads to stronger pinching on the bottom half.



(a) t =10 ns (b) t =20 ns (c) t =30 ns

(d) t =40 ns (e) t =50 ns (f) t =57 ns

Figure 5.19.: Time series of 2D axi-symmetric simulations of the solid machined X-pinch load fielded on SATURN, showing pseudo-colour of log-current density. The current is seen to diffuse into the centre, and compress on axis.

Peak compression and re-expansion

By 57 ns, the conditions in the central region are extreme. The radius drops to 25 μm from 200 μm at initialisation, giving a convergence ratio of 8. The current reaches $\sim 3/4 I_{peak}$, corresponding to $B_{\text{neck}} = \mu_0 \times (4.25 \times 10^6)/2\pi \times (25 \times 10^{-6}) = 3.4 \times 10^4$ T. The simulation is restarted at this point, with a higher temporal resolution to better observe dynamics up to peak compression. Figs 5.20(a)-(d) show the behaviour of the high pressure core over the next 800 ps. The neck has become $m = 0$ unstable and is compressed at two axial positions on either of a bulge. The upper spot reaches up to 2×10^{15} Pa at 57.6 ns, marking the point of peak compression observed given the temporal resolution. Although the simulation was run using enforced azimuthal symmetry, minor perturbations, a simplistic radiative model etc., this is still a remarkable value. The ability to reach comparable pressures to capsule implosions on the OMEGA facility [119], using cylindrical rather than spherical convergence, and on a relatively modest generator, makes this a highly interesting platform worth exploring. Details of this state are discussed in §5.3.3 for the radiation transport simulation.

Within 200 ps after this point the pressure at the hotspot drops considerably. The simulation is still far from peak current, therefore this is not due to a loss in drive pressure. Rather, this is a combination of transport of mass in the axial direction, and expansion in the radial direction. This is shown in Fig 5.21 in terms of the velocity components in the radial and axial direction. At 57.6 ns the velocity vectors are di-

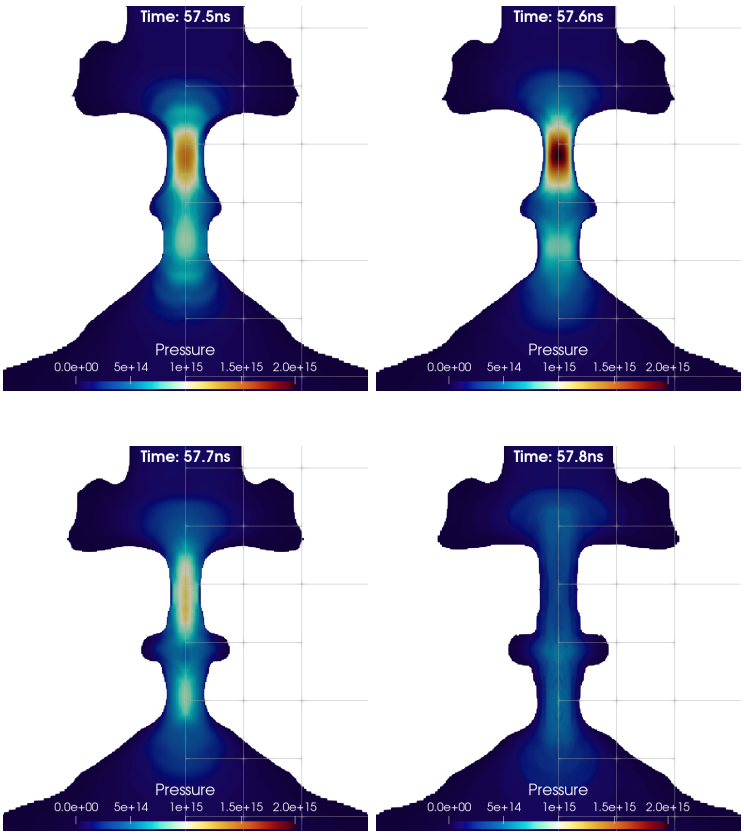


Figure 5.20.: Late time dynamics showing peak compression and re-expansion.

rected radially inwards to the hotspot but axially outwards. By the next timestep the radial components have inverted. The radial velocity magnitudes are comparable in the positive and negative times, suggesting that peak pressure occurs at some time between the two outputs.

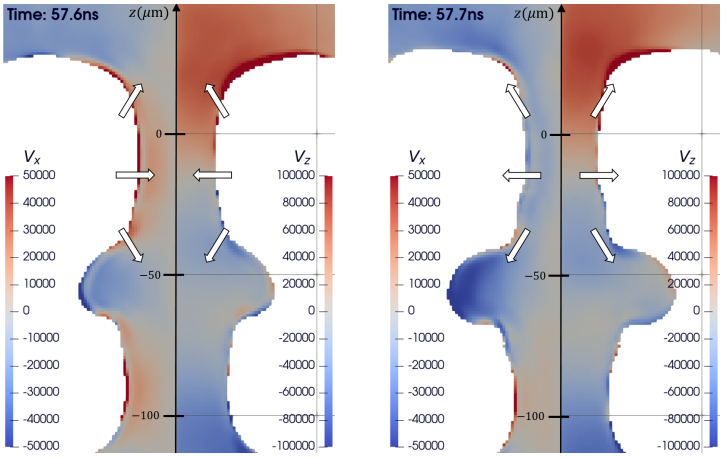


Figure 5.21.: Velocity components before and after peak compression (57.6 ns and 57.7 ns). White arrows indicate velocity direction, left half of each plot shows V_x component, right half shows V_z.

Termination

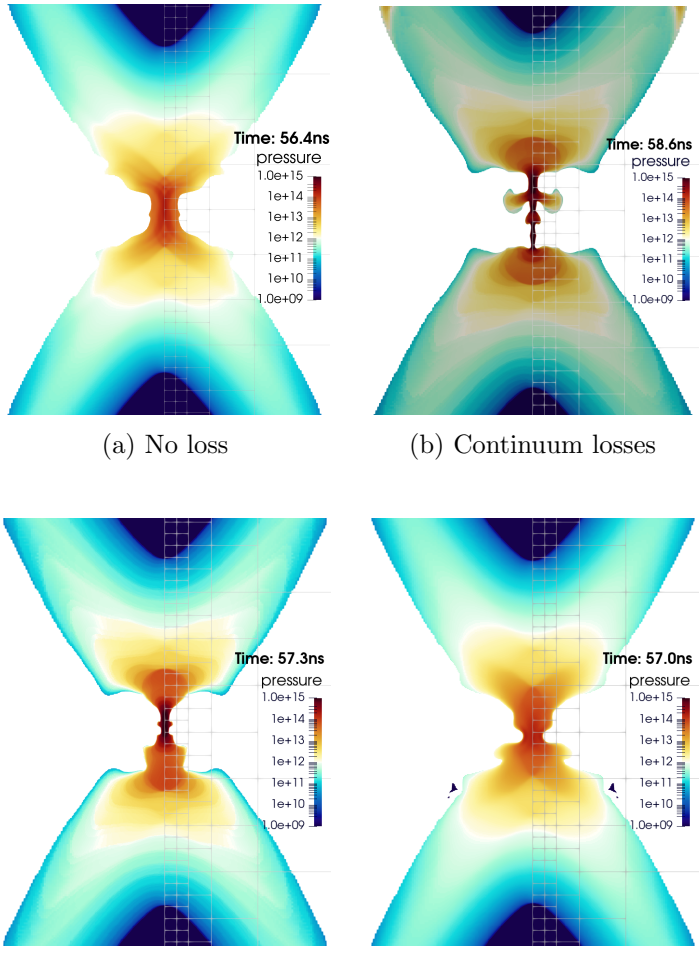
At this point, the simulation timestep drops drastically to 10^{-15} s. This is due to a combination of large magnetic field magnitudes resulting in high Alfvén speeds, and large currents that excessively Ohmically heat low density regions, generating unphysically large electron temperatures in the low density regions and requiring several thermal conduction subcycles. Strong Ohmic heating of low density plasma is frequently encountered in MHD codes, and the intense conditions of this system exacerbate this issue drastically. Continuously restarting the simulation until peak current is unfeasible for femtosecond timesteps. This early termination is unfortunate as it is likely that significant further compression occurs, since at this point only a third of the magnetic energy has been deposited, $\Delta E_{mag} = 30\%$.

Possible ways to circumvent these numerical issues include a) increasing the density floor to prevent low density material with low heat capacity to overheat, or b) applying a maximum temperature limit. The latter is undesirable as this can lead to issues with energy conservation. The former is often kept low as coronal plasma often carries significant current in wire-based systems such as wire arrays or wire X-pinches. However, it is clear in this case that the low density plumes formed here do not carry significant current (Fig 5.19(d)-(e)), and therefore are not expected to affect the broader dynamics. The only sensitivity to low-density material is how it perturbs the compression and therefore directs where along the neck most significant pinching occurs. Future simulations will be performed to confirm the insensitivity of a lower density floor, and attempt to run the simulation to completion.

5.3.3. Comparison of radiation models

The stagnation conditions described above are run with the group-resolved loss radiation model, which has been demonstrated to give good results for wire dynamics (see §4.3). The solid X-pinch, however, is a different system altogether and therefore the applicability of this model is investigated. This section compares simulations run with four different radiation model options; no-loss, simple loss, group-resolved loss, and radiation transport. Fig 5.22 shows a comparison of these four simulations close to their point of peak compression/stagnation. The numerical issues described for the group-resolved case were also encountered in the other three, which meant simulations were only run until hotspot pressures reached the first peak, $\gtrapprox 10$ ns before peak current.

The bulk plasma (electrode) region is similar between all cases, and therefore will not be discussed further. However there are considerable differences in terms of the minimum radius, and consequently the pressure, reached within the neck. The two loss cases, continuum and group-resolved, reach an evidently tighter compression as opposed to the other two. As expected, the no-loss case undergoes the least compression, reaching only 56 μm in radius, a quarter of its initial value. The electron energy losses only occur



(a) No loss (b) Continuum losses

(c) Group-resolved loss (d) Radiation transport

Figure 5.22.: Comparison of states near stagnation for different radiation models

through thermal conduction, and are insufficient to counter the expansion from Ohmic heating. This is additionally true of the coronal plasma, which becomes disruptive earlier in time.

The two loss cases undergo strong collapse, reaching the petapascal pressure range. The continuum-loss case reaches even more extreme states than the group-resolved case, close to 5 PPa (compared to 2 PPa in the latter), and collapses to 11 μm minimum radius (compared to 25 μm). In the single wire simulations, the simple loss model underestimates losses in the corona due the absence of line features in emissivity, and expands further. Here, missing line absorption features at these high densities means the simple loss contracts further, which the group-resolved loss case accounts for appropriately scales the probability of escape factor. Also in the simple loss case, the coronal plasma is free to radiatively cool and therefore does not disrupt the simulation as early. allowing it to progress by an additional 1.5-2 ns.

Further qualitative observations can be made about the structure of the neck across different models. The number of pinch spots that form is related to the degree of compression. The no-loss case, which undergoes minimal collapse, is mostly uniform. There is visible structure within the neck, indicative of shock reflections off the surrounding material that drives the radial compression. On the other hand, the radiation transport case forms a single neck that is smooth and quasi-isothermal due to radiative preheat effects (not captured in other models) diffusing shock fronts. The two loss cases go more strongly $m = 0$ unstable, particularly the simple-loss case which forms several sub-necks.

The difference in radial dynamics also has consequences in the axial direction. In all cases at earlier times, the compression from the angled electrode surface leads to a Mach stem that forms on either side of the neck. This is particularly evident in the no-loss and radiation transport cases, which steadily grows the Mach stem over time. However in the loss cases, the rapid pinch results in fast axial outflows, which pushes through the Mach stem and forms a hemispherical shock. Again, this is particularly noticeable in the continuum-loss case, where the pinching and axial motion occurs more rapidly.

Hotspot conditions

The radiation transport case represents the best available model for simulating this system, and is considered the reference case. Here, there is evidence of tight compression at one focused point, however it only reaches a minimum radius of 45 μm, a closer value to the no-loss case. The highest pressure reached in this simulation is 1.75×10^{14} Pa, a factor of greater than 10 relative to the group-resolved loss case. Fig 5.23 shows the hotspot being composed of two distinct regions, referred to here as the 'core' and 'edge' plasmas, and key parameters for each are presented in Table 5.1.

The former is shown as a high pressure central region at $1.5-1.7\times10^{14}$ Pa, which spans 19 μm in the radial direction. The latter is the high current-density region formed on the surface that spans roughly 10 μm in width. The LMD transport model predicts higher

conductivities for these near-solid, ~ 450 eV conditions, making it the preferred current path. As such, it sustains the highest magnetic field and experiences the strongest $\mathbf{j} \times \mathbf{B}$ force, providing the compression that drives the core to its high pressures. This is reflected in the difference in plasma beta parameter, $\beta \sim P_{th}/P_{mag}$. The core itself is relatively cold at < 90 eV, with the electrons, ions, and radiation temperature all in equilibrium. The difference in temperature between regions is explained by the cooling rate for a surface blackbody emitter, given by $\tau_{cool} = 0.7(Z+1)\rho h/AT^3$, where h is the height of the emitting region.

Table 5.1.: Key parameters at peak compression

Parameter	**Core**	**Surrounding**
ρ/ρ_0	8.55	0.71
Z_{eff}	17.9	32.4
P (Pa)	1.75×10^{14}	1.02×10^{14}
T_e (eV)	89.0	461
T_i (eV)	89.0	452
T_{rad} (eV)	88.9	348
τ_{cool} (ns)	159	0.17
β	0.71	$3.9{\times}10^4$

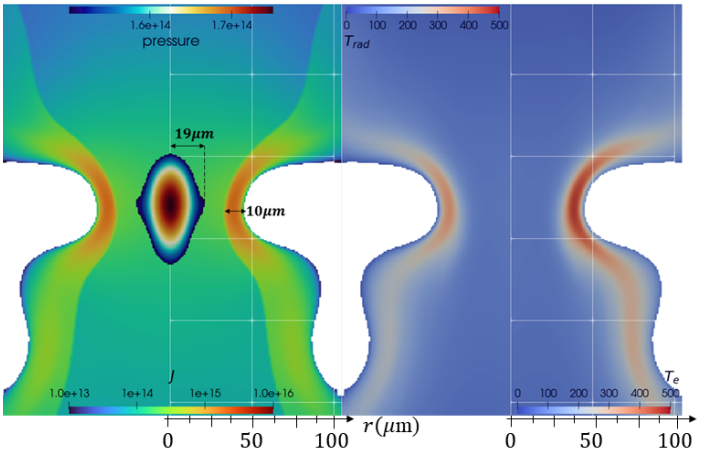


Figure 5.23.: Core and edge components of the hotspot, shown in terms of (left plot) current density and pressure, and (right plot) radiation/electron temperature on the left/right halves. The core has uniform temperatures but peaks in pressure on the axis, while the edge carries the highest current density and therefore is hotter from higher Ohmic heating rates.

5.3.4. Synthetic diagnostics

This section describes post-processing analysis of simulations to generate synthetic diagnostics for better comparison against experiments. The uncertainty associated with experimental data due to extreme environments in the target chamber is once again highlighted, making exact comparison with experiments challenging. The data used for these diagnostics come from simulations run using full radiation transport at 57 ns, where the states are described in Table 5.1, even though this does not represent the most compressed state. The processes described below represent a proof-of-principle, and therefore can be performed again once numerical issues have been addressed and simulations are run to completion.

X-ray self-emission

Synthetic self-emission images from the X-pinch at near stagnation were generated using in-house post-processing tools that solve the equation of radiative transfer along a specified line-of-sight (LOS):

$$I_\nu(s) = I_\nu(s_0)e^{-\tau_\nu} + \int_{s_0}^{s} j_\nu(s')e^{-\tau_\nu}ds', \tag{5.13}$$

where $\tau_\nu = \kappa_\nu ds$ is the frequency-dependent optical depth, and s is the distance propagated by the radiation [74] [75]. For these simulations, the simulation grid is first interpolated onto a high resolution uniform grid, as the ray tracing tool is not yet compatible with variable resolution grids.

Self-emission images spatially integrate information along an LOS, making it challenging to decipher the composition of the source. An additional complexity arises from the low spatial resolution and 35° viewing angle in experimental images. It is therefore instructive first to compare synthetic images from a side-on view with simulation output, as shown in Fig 5.24. The density and electron temperature profiles used are shown in the top two figures, while the bottom two show results of post-processing using 250±45 eV and $>$ 1 keV spectral filters for soft and hard emission respectively. The logarithm of intensity is used which more clearly reveals less bright features for better comparison against simulation.

The most obvious feature in both soft and hard emission is the central hotspot, which appears as a $r \sim 50$ μm brightly emitting region, corresponding to the radius of the edge rather than core region described above. No structure is visible within the bright spot at the core radius. Furthermore, the photon mean free paths λ_{mfp} across almost the entire spectral range at the edge and core conditions lie below the 25 μm radius of the hotspot (Fig 5.25(a)). Therefore emission will be from a surface blackbody at edge temperatures, $\sim$ 450 eV.

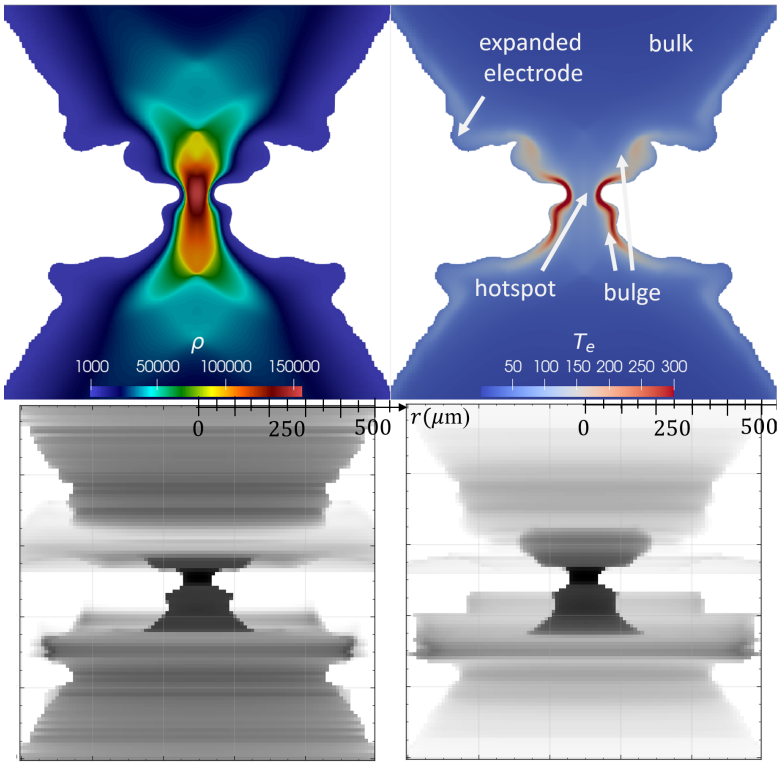
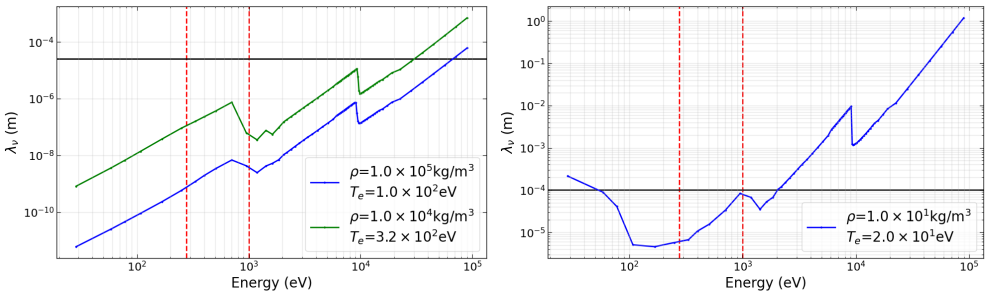


Figure 5.24.: *Top:* density (left) and electron temperature (right) profiles from simulation run with radiation transport at 57 ns. Synthetic self-emission images using 290 eV (left) and > 1 keV (right) filters, from a side-on LOS. Horizontal striations are due to 2D simulation output being azimuthally revolved to imitate 3D output.



(a) λ_{mfp} for the core (b) λ_{mfp} for the expanded electrode

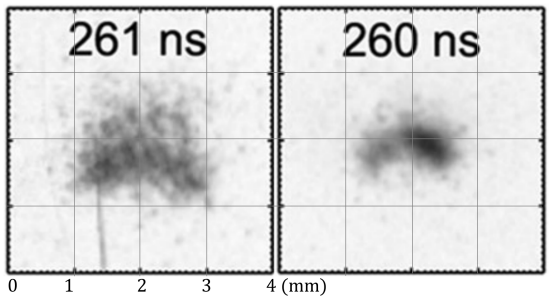
Figure 5.25.: Photon mean free paths $\lambda_{mfp} = 1/\kappa_\nu$ as a function of energy for two regions in the X-pinch.

Further away from the neck in the axial direction, the next brightest emitters are the two bulges on opposite ends of the hotspot that extend to 90-200 μm radius, appearing in the image in dark grey. These are significantly less dense than the hotspot (0.5× solid) but still optically thick for most of the spectrum, E <10 keV. However, in the soft emission image the top half is occluded by low density emission from the expansion of 'electrode' material, a feature which is less visible in the hard emission. Once again looking at λ_{mfp} for this feature (Fig 5.25(b)) it can be seen that the 277 eV radiation

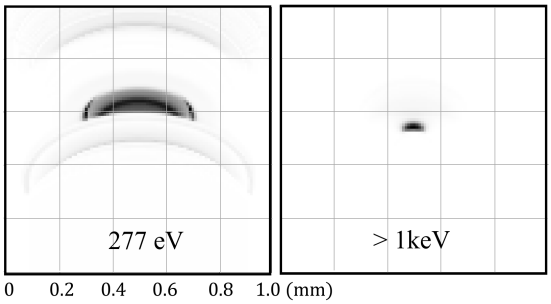
is optically thick ($\tau \sim 0.1$) while the >1 keV emission is optically thin.

With this structure from a side-on view in mind, the images from a 35° LOS are now discussed. Fig 5.26(a) shows the experimental image of approximately the same time in the current pulse (60 ns in experiment vs 57 ns in simulations), and Fig 5.26(b) shows the corresponding synthetic images. Note, detector response to emission varies linearly, so the linear rather than logarithmic data is shown. Also, the spatial scales of the two images are different; 4×4 mm^2 for experimental and 1×1 mm^2 for synthetic images.

In both sets of images the brightest feature has distinct curvature due to expanded material from the lower electrode blocking emission. In the experimental image it appears as tall as it is wide, roughly 3×3 mm^2. Here, the hard emission coincides with the bright subset of the soft emission band, with a localised, high-intensity spot towards the right. In the synthetic image, the dominant soft emission is a narrow band extending 400 μm, corresponding to the top bulge above the hotspot. Like in experiment the synthetic hard emission is localised, however it is located directly below the soft emission band. This originates from the hotspot itself, which is obscured in the lower energy range by the expanded electrode material.



(a) Experimental



(b) Synthetic

Figure 5.26.: Self-emission images from a 35° LOS, showing (a) experimental images and (b) synthetic images, for (left) lower and (right) high energy emission.

X-ray streak camera

§5.3.1 describes the use of an x-ray streak camera with several filters to get an estimate of the emission spectrum at pinch-time. These are compared against synthetic x-ray intensities using two assumed sources: a 1275 eV blackbody and 3000 eV bremsstrahlung. An equivalent process can be performed using the spectrum from simulations.

To ensure consistency, this process is first reproduced for the blackbody spectrum, shown here in Fig 5.27(a). The left plot shows digitised data for the filter response $f(\nu)$ (which includes the streak camera response contribution) for four filters; Sc, V, Ti, and Zn, overlaid with the blackbody spectrum, $B_\nu \sim \nu^3/e^{-h\nu/k_BT}$. The right plot in Fig 5.27 shows the cumulative sum of the product of filter response (f_ν) and spectrum (B_ν):

$$C_i = \int f(\nu)B_\nu dE \simeq \Sigma_{\nu_i} f(\nu_i)B_{\nu_i}\Delta E. \tag{5.14}$$

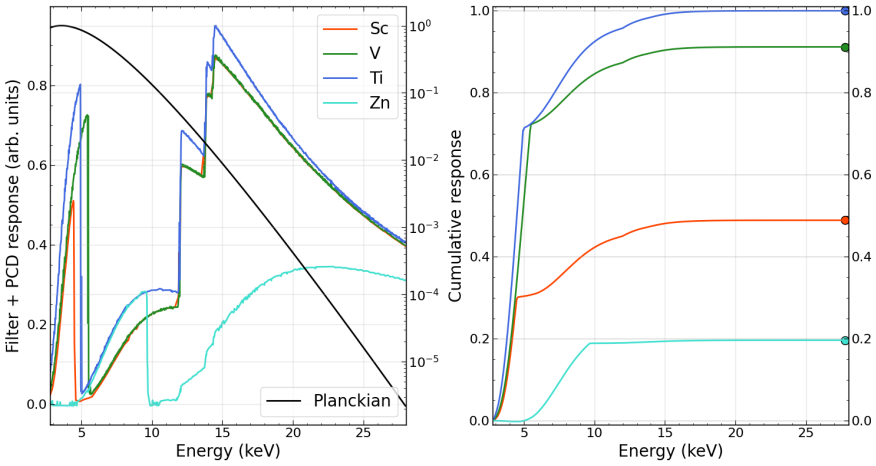
By interpolating the data onto a uniform energy stencil the factor dE can be neglected. The ratio of the cumulative sums for each filter material are presented in Table 5.2 in the "B_ν" column, in brackets. The values not in brackets are the original calculations presented in Ref [22]. There are small differences between original and reproduced values, however these are attributed to errors in filtered data from digitisation.

The above process is then performed using the spectrum obtained from simulation. The simulation was restarted from 57 ns, but using tables with high spectral resolution, i.e. 400 energy groups instead of 32. This number of groups would have been unfeasible to run for the full simulation, however here the restart was only performed for ~ 50 ps until the radiation transport calculation settled to a steady state. A probe was then placed in the vacuum region away from the X-pinch to output the energy density as function of frequency group. The resultant spectrum is shown in Fig 5.27(b), again overlaid on the measured filter response with the cumulative sums on the right plot. A distinctive recombination edge feature is visible in the simulation spectrum at 11.5 keV, corresponding to the copper K-shell edge [120]. This feature is an artifact of using scaled copper atomic data as a surrogate for tungsten, and it is expected that using correct tungsten data would more closely produce a blackbody spectrum, with a low-intensity L-shell feature at 12.1 keV [120]. The ratio of cumulative sums are presented in the "Simulation" column of Table 5.2.

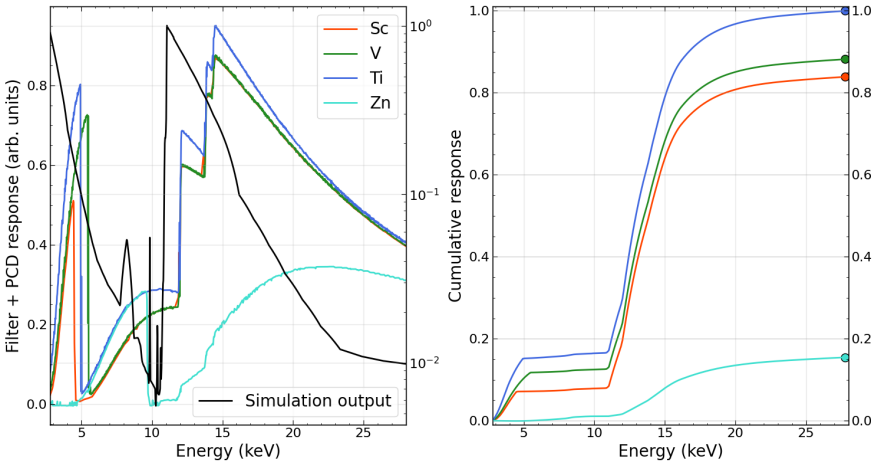
The set of ratios obtained from the blackbody and simulation spectra can be compared against the experimentally obtained amplitude ratios, shown in Fig 5.16(c). The blackbody spectrum shows reasonably good agreement with experimental values, while the simulation ratios are considerably different, particularly for the Sc filter. This can be explained using the two right-hand plots of Figs 5.27(a) and (b). For the 1275 eV blackbody spectrum, the cumulative sum diverges between filters in the 3-5 keV range where the spectrum peaks, which weights the ratios towards the difference in filter amplitudes at this energy. The difference across the remaining spectral range is relatively constant. However for the simulation spectrum, the lines diverge at the 11.5 keV recombination edge, which weights the cumulative sum towards the filter ratios at this point, giving an unphysically high contribution of the scandium filter. This suggests the true spectrum in experiment would have been weighted in the lower, 3-5 keV, range, and that using Cu-scaled tables is inappropriate for direct comparison against spectral diagnostics for a tungsten plasma.

Table 5.2.: Streak signal amplitude comparison

Filter	Experimental	B_ν (reproduced)	Simulation
Sc	0.625	0.50 (0.49)	0.84
V	1.0	0.88 (0.91)	0.88
Ti	1.0	1.0 (1.0)	1.0
Zn	0.21	0.22 (0.20)	0.15



(a) 1275eV blackbody spectrum



(b) Simulation spectrum

Figure 5.27.: *Left:* digisted filter responses in coloured lines, spectrum used for calculation in black line. *Right:* cumulative sum as function of photon energy using Eq 5.14 for different filters. Note, all values shown here are normalised to the maximum value within each data set, as absolute measurements are not available.

5.4. Concluding remarks

The chapter presents simulations of experiments performed using two configurations of X-pinches: wire (WXP) and solid (SXP) X-pinches. Both configurations have demonstrated the formation of extreme conditions in their hotspots due to strong compression of the crossing point.

The hotspot pressures reached in wire X-pinch case (4TPa) was found to be considerably higher than the hotspots formed along the single wire from the m=0 mode (0.2GPa). This was despite comparable currents delivered to each of the loads within a similar rise time ($\dot{I} = 2.3 \times 10^{12}$ A/s for IMP versus $\dot{I} = 3.6 \times 10^{12}$ A/s from XP to the target X-pinch), and a lower resolution used in the WXP case. However, the WXP simulations were terminating before further compression could occur. An obvious next step is to increase the domain size and run the simulation to completion to determine whether even higher pressures are possible. Additionally, development of a synthetic radiograph capability is current underway for 3D simulation output using mesh refinement. Comparison of experimental versus synthetic radiography of the compressed state is an objective for further studies of this system.

The WXP parameters used in simulation were based on which parameters had the most extensive experimental dataset, for comparison. These were Mo X-pinches, which has a time series of radiographs showing the evolution of the WXP up to 200 ps before its pinch event. However, other diagnostics were not fielded in the Mo experiments, such as the step wedges used in Al experiments for areal density estimates. Quantitative comparisons of, for example, the areal density of the axial jet could be used to determine whether the ablation off the wire surfaces is being correctly modelled. Therefore, a more extensive parameter scan using other materials would be an informative avenue for further investigation. Finally, simulation 2-wire X-pinch exhibited significant 3D asymmetry in the overall dynamics due to the offset of wires from the centre, although the final compression at the neck was uniform. It would be interesting to obtain experimental data in the plane perpendicular to the X-pinch rather than the standard side-on view to confirm this asymmetry.

The simulations in §5.3 are the first detailed study of the SXP system performed on the SATURN generator. The extreme magnetic energy deposited into the central crossing point make this a highly interesting system. Simulations presented above were performed up to the first peak in pressure, however the X-pinch neck remained suggesting further compression is possible at later times. A study of this system using different radiation treatments were performed. Simulations performed using the simple loss model (accounting only for continuum processes) was shown to undergo radiative collapse to $\sim$10 μm radius. Using radiation transport instead showed a termination of the collapse due to absorption processes occurring in the neck as it become optically thick. Future studies will look to investigate if further compression occurs at later time using the higher fidelity radiation model, and whether the 5 μm spot size inferred in experiment can be attained.

6. Magnetic reconnection

This chapter studies radiative cooling in an entirely different system to the previous two chapters, namely in the reconnection layer formed between two adjacent exploding wire arrays. Simulations performed in 2D do not capture the effects of axial modulations along the wire surface, which can potentially disrupt plasmoid formation in the layer. The layer formed in this system is a thin, largely stationary region in the centre of a large domain, necessitating a mesh refinement capability when simulated at high resolution in three dimensions.

6.1. Overview

Magnetic reconnection occurs when two anti-parallel field lines are brought together in a central layer and undergo a change of topology, resulting in a large release of stored magnetic energy. The curvature of the rearranged field lines generates magnetic tension, producing outflows that move at the order of the Alfvén velocity, as shown in Fig 6.1. The central shaded region where field lines break and reconnect is referred to as the 'reconnection layer', with a length L and width δ. In steady-state, the reconnection 'rate' (often defined as a dimensionless parameter) is of the order of the ratio of inflow velocity to outflow velocity, V_{in}/V_A.

This process is ubiquitous in space and astrophysical plasmas, and is thought to occur in extreme environments such as black hole accretion disks and pulsar magnetospheres [121]. More terrestrially, reconnection occurs on the sun-side of the Earth as its magnetosphere interacts with the solar wind, causing field lines to be deflected around its

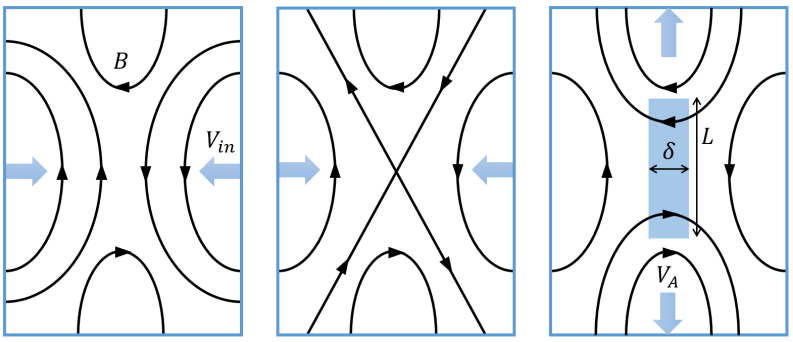


Figure 6.1.: Diagram of reconnection process, showing magnetic field lines in black and velocity arrows in blue. From left to right: field lines are advected into the central layer at the inflow velocity, V_{in}, undergo a rapid change of topology, and are advected out in the perpendicular direction at Alfvèn velocity, V_A.

surface and undergo reconnection again in the magnetotail [122]. Despite its frequency of occurrence in the universe, several unanswered questions remain about this process. For example, observed reconnection rates are often significantly faster than predicted by existing simple models.

The Sweet-Parker model was developed separately by Sweet [123] and Parker [124], and is a simple theory often used to describe a collisional, resistive reconnection layer. Consider the system in Fig 6.1, which forms a laminar current sheet at the layer. Conservation of mass requires that the rate of mass entering the reconnection layer is equal to that exiting it:

$$\rho V_{in} L = \rho V_A \delta. \tag{6.1}$$

The electric field due to the inflow entering the layer can be described by Ohm's law for ideal MHD, $E_{in} = V_{in} B_0$, where B_0 is the magnitude of the field carried into the layer. Assuming a resistively dominated layer, the electric field in the outflow can be described by $E_{out} = \eta j$. Equating electric fields and using Ampere's law in a plasma, $j \sim B_0/\mu_0 \delta$, yields an expression for the width of the layer:

$$\delta = \frac{\eta}{\mu_0 V_i n}. \tag{6.2}$$

Eq 6.1 can be rearranged to give the steady-state reconnection rate $V_{in}/V_{out} = \delta/L = S^{-1/2}$, where S is the Lundquist number, which is a measure of the plasma Alfvén transit time versus resistive diffusive timescale, and is defined as:

$$S = \frac{\mu_0 L V_A}{\eta}. \tag{6.3}$$

This model, while used for several decades, was found to only match the observed reconnection rate in high resistivity, low S regimes. In recent years a description of continuous, bursty ejection of magnetic flux has emerged [125]. MHD simulations repeatedly showed the central current sheet becoming unstable to the 'plasmoid' instability, where a chain of magnetic islands ('O points') called plasmoids form along the layer (as shown in Fig 6.2) and are rapidly ejected with the outflows [24] [23]. This picture leads to a higher rate of magnetic flux being exhausted from the system, and therefore a higher reconnection rate as compared to the Sweet-Parker model.

The transition between the first steady-state description and the second spurious phenomenon is shown in Fig 6.3 as a function of Lundquist number and the layer length normalised to the ion skin depth, $\lambda = L/d_i$ [24]. These are referred to in the diagram as 'single X-line' and 'multiple X-line' regimes respectively, and are separated by transition criteria, $\lambda = \lambda_c$ and $S = S_c$. The first criteria is used in the limit of infrequent collisions or large scale length relative to mean free paths (higher S lower λ). The second is used in the case of collisional plasmas (lower S higher λ), which is of relevance to the system described in this chapter. Here, the transition to a plasmoid instability is shown to occur at $S = S_c \sim 10^4$ as determined by MHD simulations [24].

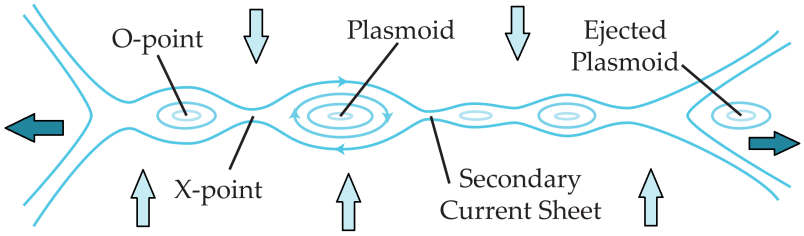


Figure 6.2.: Diagram showing plasmoid formation at the centre of a reconnection layer, and ejection with the outflows. These are characterised by 'O' point topology of the magnetic field. Lightly shaded arrows indicate direction of inflow, darker arrows indicate outflows. Adapted from Ref [23].

The above discussion has so far only considered MHD effects; in reality several processes occur in reconnecting systems. For example, extreme astrophysical environments are often radiation dominated, which influences the inflow and outflow states of the layer. Udzensky and McKinney [126] provided the first treatment of radiative effects on the Sweet-Parker layer described above. In this model the Ohmically heated reconnection layer additionally dissipates thermal energy via optically thin radiation. Assuming Spitzer resistivity and continuum-only (bremsstrahlung) emission, the layer width can then be shown to depend on:

$$\delta^2 \sim \frac{\ln \Lambda}{n_0} \left(\frac{V_{A0}}{c} \right)^{-2}, \tag{6.4}$$

where V_{A0} and n_0 are the upstream Alfvén velocity and plasma density respectively. Several experiments have been conducted to study reconnection across various regimes.

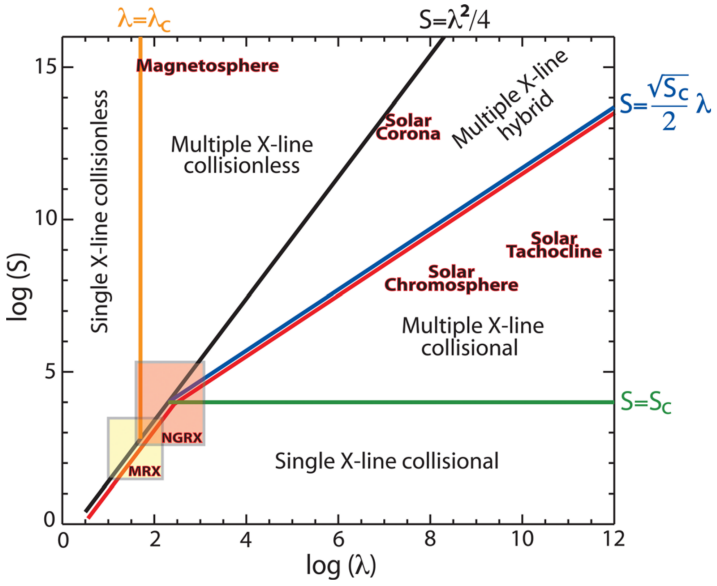


Figure 6.3.: Phase diagram of magnetic reconnection regimes, in terms of normalised scale length λ and Lundquist number S. Reproduced from Ref [24].

In the low-density, large magnetic Reynolds number limit is the MRX (Magnetic Reconnection eXperiment) platform, which applies external magnetic fields to a large gas-filled vessel [127]. In the opposite limit, laser-driven experiments have been conducted using two beams to irradiate adjacent regions of a solid target, forming bubbles that produce anti-parallel self-generated magnetic fields from the Biermann battery effect [128]. The first of these experiments is easily controlled and well diagnosed, the second is collisional but extremely transient ($\sim$ 1ns) and therefore challenging to probe.

Ref [129] describes the development of pulsed-power driven platform to study collisional reconnection on the MAGPIE facility at Imperial College London, over a long timescales (hundred of nanoseconds) and large length scales ($\sim$mm). This system uses two adjacent wire arrays in the exploding configuration, each of which carries azimuthal magnetic fields radially outwards. The region where they collide forms a relatively uniform reconnection layer, which can be studied using MAGPIE's extensive diagnostic suite. The platform allows for flexible modification of design parameters, for example the inflow density depends on the current per wire, which itself is controlled by the number of wires per array. This configuration was later fielded on the Z facility at Sandia National Laboratories as part of the MARZ (MAgnetical Reconnection on Z) campaign and experiments are underway at the time of writing. These recent experiments are the focus of the simulation studies conducted in this chapter.

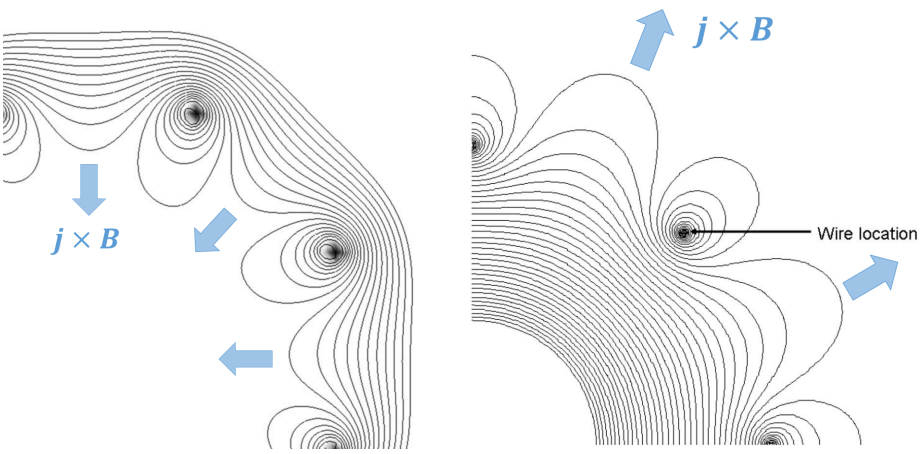
The remainder of this chapter is organised as follows. A description of wire array dynamics, in the imploding and exploding configurations, is first provided in §6.2. Also included is a review of exploding array reconnection experiments performed on the MAGPIE facility, and a description of the modifications required to field this configuration on Z. §6.3 describes numerical investigations using 2D simulations of the effects of resolution and radiative models on this system. Finally, §6.4 describes the first 3D simulations of this system to resolve the reconnection layer at 25 μm, enabled by the mesh refinement capability.

6.2. Exploding wire arrays for magnetic reconnection

In their standard configuration wire arrays are typically used as high efficiency soft x-ray sources, however the ablation streams are challenging to probe along the axial direction as they move radially inward and are diagnostically blocked by the wire cores. In the inverse configuration, the ablation stream diverge radially while maintaining similar properties as the standard configuration, making them useful in a variety of applications. This section describes the behaviour of imploding arrays, exploding arrays, and their application to 1MA reconnection experiments.

6.2.1. Imploding arrays

The single wire system of §4 is subjected to a (mostly) radial Lorentz force due to its axial current and local azimuthal magnetic field. In the imploding configuration,



(a) Imploding array (b) Exploding array

Figure 6.4.: Magnetic field topology for a standard imploding wire array (left) and exploding wire array (right), both of which have axially directed currents. Wire cores in each configuration have a local field which merge with the global field. Direction of resultant $\mathbf{j} \times \mathbf{B}$ force is indicated by blue arrows. Modified from [25] and [26] respectively.

wires are arranged in a uniformly-spaced circular array. The current path is closed using return posts, similarly arranged in a circular array but at a larger radius. The global magnetic field formed outside the wires is nearly azimuthal as shown in Fig 6.4(a). Therefore, the resultant $\mathbf{j} \times \mathbf{B}$ force is directed radially inwards towards the array axis, as opposed to the wire axis as in the single wire case. The ablation streams from each wire, which are accelerated to $50 - 100$ km/s (on MAGPIE), undergo a series of oblique shock interactions with streams from adjacent wires before they collide on axis. Once sufficient mass has ablated off the wire surface, the wire cores begin to break along the region of axial modulations and undergo a 'snowplough' implosion. The wire cores do not uniformly get swept axis-ward, rather they are susceptible to MRT (magneto-Rayleigh Taylor) instabilities that form on the outer wire surface. These MRT instabilities arise when radially ablated coronal material on the outer wire surfaces are redirected around the core and towards the axis. The magnetic field behaves as a light fluid pushing against the corona, non-linearly exacerbating perturbations along its surface.

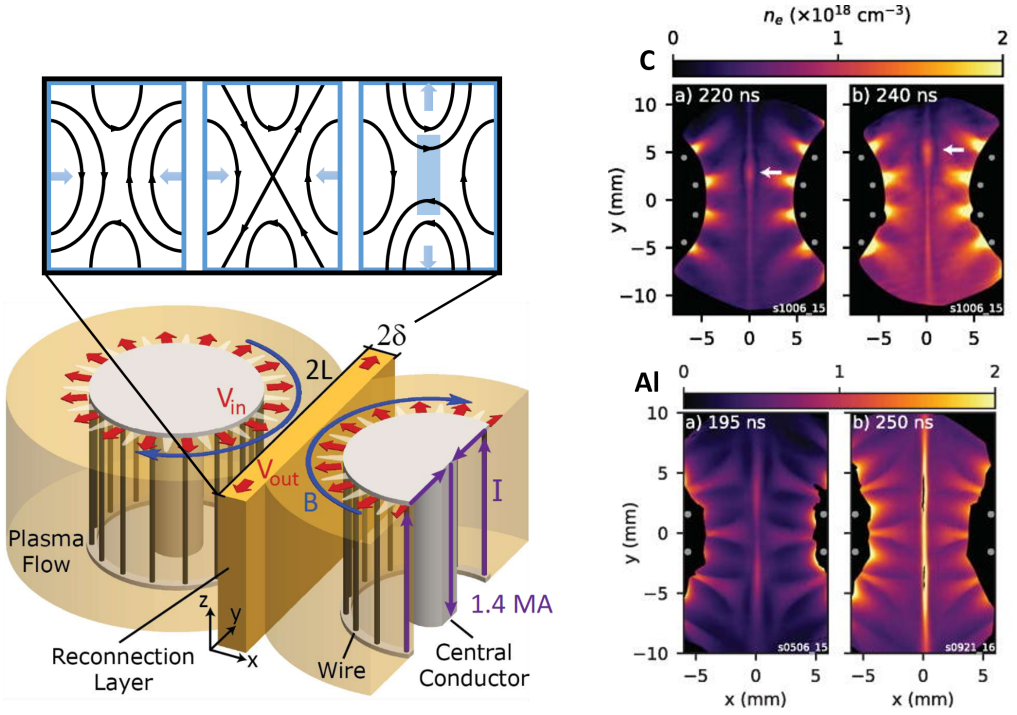
6.2.2. Exploding arrays

The behaviour wire arrays in the inverse configuration is similar to the standard configurations, except for certain dynamics that occur in the oppositely oriented direction. Firstly, the ablation streams diverge radially outwards rather than converge on axis. The global magnetic field is formed around the central cathode rod instead, reversing the direction of the resultant $\mathbf{j} \times \mathbf{B}$ force as shown in Fig 6.4(b). Another difference is the surface along the wire that becomes MRT unstable; in the exploding configuration it is the inner surface where instabilities are seeded. This is discussed further in §6.4. However, the magnitude of the force experienced by each wire is similar to the equivalent

imploding array with the same array parameters. Also similar to the imploding case the expanded flows around the wires undergo shock reflections with adjacent streams. As such, the radially diverging outflows have previously been used for several purposes including to study the properties of wire ablation from imploding arrays [26], to study magnetised bow shocks around obstacles [130], or as a current switch to drive faster implosions of a standard array [131].

6.2.3. Reconnection on MAGPIE

This section describes results from Ref [129], which describes the first set of published experiments using the pulsed-power driven reconnection platform described above. By default these used two exploding wire arrays with 16mm diameter and made of 16 wires each, separated centre-to-centre by 27mm. A diagram of this set-up is shown in Fig 6.5(a), where the ablated plasma is shown in transparent yellow, and the interaction (reconnection) layer is shown in solid yellow. Diagnostics used to probe the inflow plasma and reconnection layer were primarily laser interferometry for line-integrated electron density maps [27], and optical Thomson scattering to obtain flow velocity, electron temperature, and ion temperature measurements [132].



(a) Experimental set-up (b) Density maps

Figure 6.5.: Set up and results from experiments conducted on the MAGPIE facility. *Left:* diagram of two adjacent exploding wire arrays, showing direction of magnetic field and velocity. Illustration of reconnection processor highlighted along layer. *Right:* processed interferogram images shown electron density maps for carbon (top) and aluminium (bottom) experiments. Modified from Ref [27].

Array parameters were varied to change conditions in the inflow, for example the array separation was increased/decreased to control the plasma density and magnetic field magnitude transported into the layer, and the planarity of the inflow which determines the layer length L. Of relevance to this study, however, is the change in dynamics between different wire materials. Two sets of experiments were conducted, using carbon (Z=6) and aluminium (Z=13) wires. Table I in Ref [129] provides a list of several plasma parameters in the inflow [1] and layer regions separately, and for each of the materials used, two key differences are described here.

- While both wire materials had inflow velocities of $v \sim 50$ km/s, the plasma beta parameter $\beta \equiv P_{th}/P_{mag}$ was calculated to be approximately 1 in the aluminium case and 0.4 in the carbon case. Both materials were found to have similar thermal pressures, therefore this was due to higher magnetic fields in the carbon case, leading to a stronger contribution of magnetic versus thermal pressure.

- In both regions, and for both materials, the ion temperature was observed to be higher than electron temperature. However within the layer, the rate of electron-ion collisions was calculated to be larger for the higher-Z case, implying that electrons were heated from the thermalised ion kinetic energy more efficiently. This dropped the radiative cooling timescale, which was calculated to be $\tau_{rad} \approx 4$ ns for aluminium and 600 ns for carbon.

The consequence of these differences is a difference in reconnection dynamics in the layer. Electron density maps obtained from top-down interferogram images at two different times for both materials are shown in Fig 6.5(b), showing visibly different structure. The carbon images show the presence of a high density elliptical feature moving out of the layer (marked by the white arrow), which can be identified as a plasmoid. This is further supported by magnetic field measurements from B-dot probes positioned at the layer outflow, which show a sinusoidal profile in amplitude indicative of a magnetic 'O' structure. In contrast, the aluminium layer appears uniform along its length, and continuously compresses to increasingly higher densities. From the parameters above, this behaviour can be attributed to strong radiative cooling which drops the thermal pressure and allows a pile-up of density (and magnetic field) along boundary. The cooling further increases the resistivity ($\eta \propto T^{-3/2}$), which drops the Lundquist number in this region and pushes it out of the 'multi X-line' (plasmoid forming) regime in Fig 6.3.

6.2.4. Reconnection on ZR

The 1 MA experiments were observed to rapidly undergo a transition from the plasmoid-unstable regime to the radiatively cooled, resistively-driven regime, making it challenging to observe the formation of the cooled reconnection layer. The larger, 20MA Z facility allows access to regimes of higher Lundquist numbers due to the size of the

[1] For the remainder of this chapter the terms 'inflow' and 'outflow' are used with respect to the layer, therefore the layer 'inflow' is the array outflow

arrays and reconnection layer formed, as well as higher radiative cooling rates from the larger currents driving the reconnection layer to higher densities. The MARZ experiments performed on Z attempt to achieve these conditions using two 40mm diameter arrays each comprising 75 or 150 (100 μm) wires and separated by 60 mm [77].

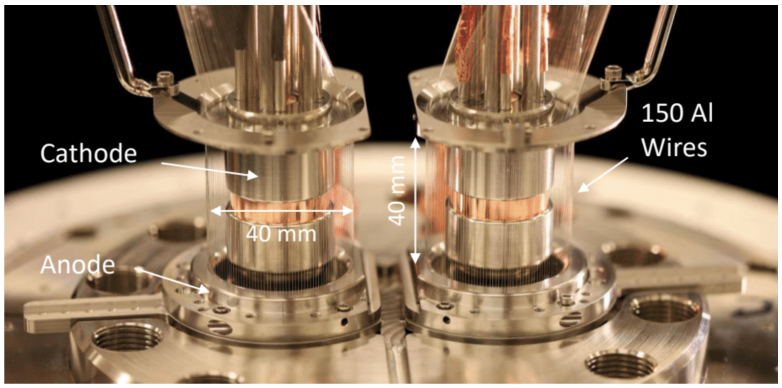


Figure 6.6.: Image of MARZ experiment load hardware, showing two wire arrays in exploding configuration with relevant dimensions labelled. Courtesy of Sandia National Laboratories.

6.3. 2D Simulations

Simulation studies performed in 2D of the above system are described here, with a focus on the effects of resolution and radiative models. A comprehensive study of the effect physics models, grid configurations, array parameters, etc. are beyond the scope of this work. The versatility of this platform, while useful for controlling conditions in the array, makes it challenging to disentangle the complex interplay of each effect. Instead, two specific comparisons are made. The first is the effects of resolving the reconnection layer separately to the wire dynamics, which can only be achieved using a variable mesh capability. The second is the effect of radiation model on the layer dynamics, to complement the study performed in Ref [77].

6.3.1. Set up

By default the simulations presented here were run using 75 wires in each array, and at $\Delta x = \Delta y = 50$ μm resolution over a domain size of $L_x \times L_y = 16 \times 8$ mm^2. This cell size was determined to be smaller than the reconnection layer resistive diffusion length scale $\eta_{\text{layer}}/V \sim 200$ μm, and therefore sufficient to resolve the dynamics [77]. Also at this resolution, the wires were initialised as 3×3 cells squares to ensure symmetric ablation in the $\hat{x}$ and $\hat{y}$ direction. This cross-sectional area overestimates the linear mass using solid density, therefore the initial density was scaled down by a factor of $f = \pi r_w^2/9\Delta x \Delta y$ to give $\rho_0 = f\rho_s = 942$ kg/m^3. Temperature profiles were initialised at 0.1 eV with a %1 perturbation, as described in §4 for 2D and 3D single wire simulations.

Default FEOS and LMD models were used for equation of state and transport respec-

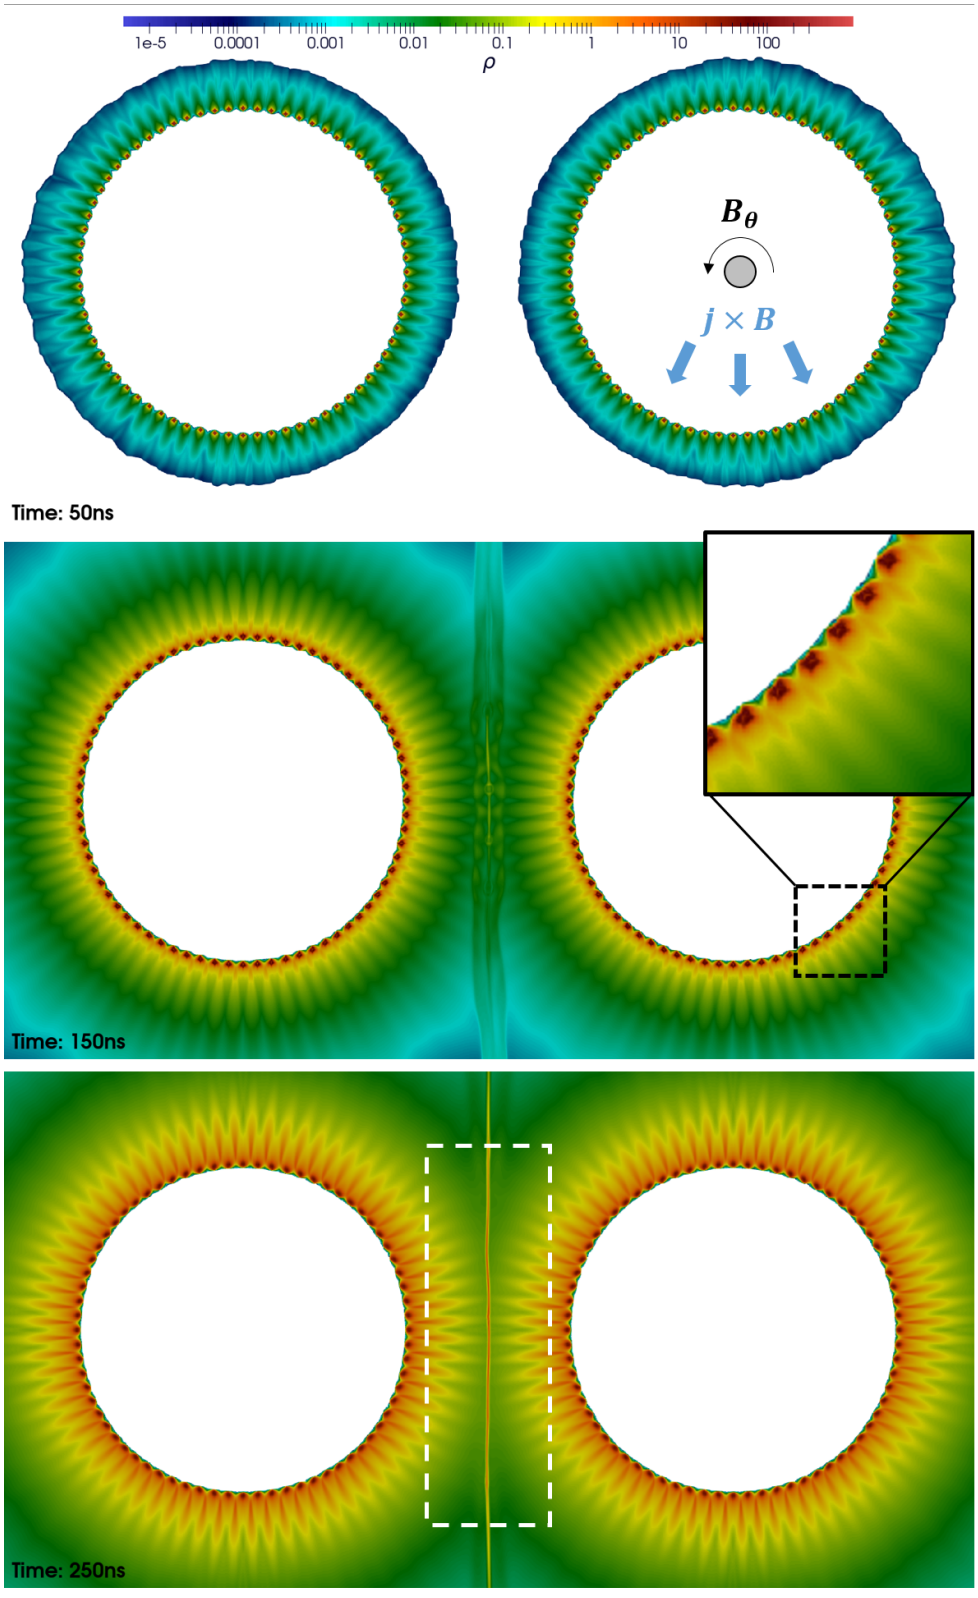


Figure 6.7.: Evolution of simulations of the MARZ experiment, shown at three different times (50, 150, and 250 ns from top to bottom) in terms of the logarithm of density in kg/m^3. Top image right array shows position of central conducting rod relative to array.

tively. Additionally, the group-resolved radiation loss model was used for simulations presented in §6.3.2 and 6.3.3, and comparisons against a transport model is described further in §6.3.4. The current pulse in experiment is approximated by a $\sin^2$ profile, with $I_{peak} = 20$ MA and $\tau_{rise} = 300$ ns, corresponding to the parameters of Z when operated in the long-pulse mode [133].

6.3.2. Dynamics

The dynamics of a typical simulation are demonstrated in Fig 6.7, showing the density profiles at three different times. These are described first in terms of the array outflows then in terms of the layer (white dashed box in bottom image). The first plot at 50 ns shows the ablation streams ($\rho \leq 0.1$ kg/m^3, green to blue contours) emerging from the wire cores (red dots). Also annotated on the right array is the position of the central conductor in a grey circle, and the direction of the magnetic field and Lorentz force in black and light blue arrows respectively. The outer vacuum-plasma boundary is almost uniformly circular; observed variations are due to the applied temperature perturbations. The discrete nature of the ablation is most visible close to the wire cores. The interaction between individual streams is shown more clearly in the inset at 150 ns. At each wire, the radial ablation from the inner surface is rapidly redirected around the wire and accelerated outward. Each redirected stream reflects off its neighbour, causing a local density increase and raising the ion temperature through viscous heating. At this time, the flows have been accelerated up to 120 km/s before being decelerated by the stagnated layer. By 250 ns the current pulse has almost reached its peak. The wire cores continuously replenish the supply of ablated coronal plasma, although at increasing density to match the rising current. However, the most significant change is in the reconnection layer, described next.

Reconnection layer

A magnified view of the layer itself is shown for the two later times, 150 ns and 250 ns, and across different variables, in Fig 6.8. The structure observed at 150 ns in the interaction region on either side of the narrow density peak is shown more clearly in the density profile (third panel). The peaks and troughs coincide with interaction regions between ablation streams, and are the result of a reflected shock off the axis. The magnetic field also experience a pile-up on axis as seen on the leftmost panel, where the ratio of magnitudes in the shocked versus unshocked region is $B_s/B_0 \sim 3$. This is additionally reflected in the current density profile, which has heightened values along the layer surface. Most importantly, multiple plasmoids are formed along the layer. These show signatures in all four variables; 'O' shaped magnetic fields, enhanced current and mass densities, and elevated electron temperatures. Furthermore, the top and bottom plasmoids are rapidly ejected along the layer out of the simulation domain at supersonic, super-Alfvénic velocities ($M_s = 4.5, M_A \sim 2.6$). However, between $170-190$ ns, the electron temperature in the central two plasmoids rapidly drops as the density

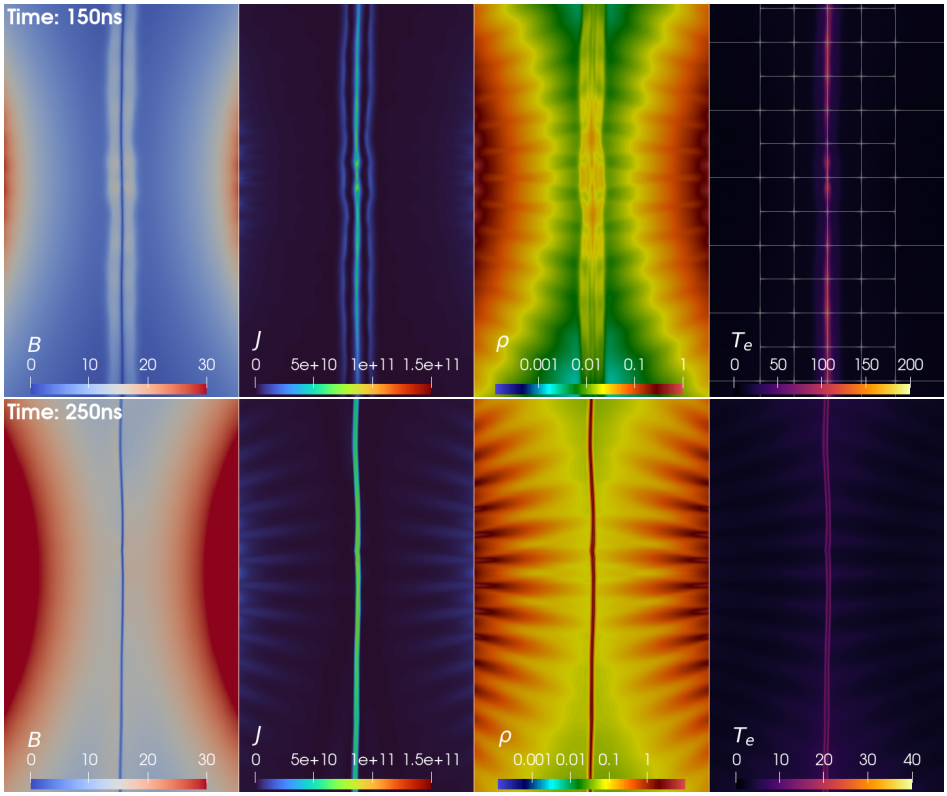


Figure 6.8.: Evolution of the reconnection layer at 150 ns and 250 ns for (from left to right) magnetic field in T, current density in A/m^2, density in kg/m^3, and electron temperature in eV.

continues to increase. The brightly emitting regions steadily decrease in width until they are extinguished entirely by 250 ns, as shown in the bottom series of panels. Only the central current sheet remains, which gets increasingly compressed to higher densities, further increasing the cooling rate and dropping the temperature. The resistivity in this region increases from its value at 150 ns by a factor of $5 - 6\times$, pushing the layer further into a low Lundquist number, single 'X-line' region.

6.3.3. Resolution effects

The resolution used above has been demonstrated both analytically and through convergence studies to be sufficient to drive the system into an initial high Lundquist number, plasmoid-generating regime, and later model the collapse of the layer to resistivity-driven reconnection. The layer reaches a minimum thickness of 200 μm through radiative cooling, and is therefore typically resolved across 4-5 cells. Results presented in Ref [77] were similarly run at $\Delta x = 50$ μm, however simulations are increasingly being performed at 25 μm. Key differences in behaviour between these resolutions are firstly more plasmoids being formed along the layer, and secondly a more rapid rate of extinguishing. These are often attributed to a difference in layer Lundquist number due to lower numerical resistivity, however the observations cannot be decoupled from the effects of wire ablation resolution. A mesh refinement capability offers the flexibility to

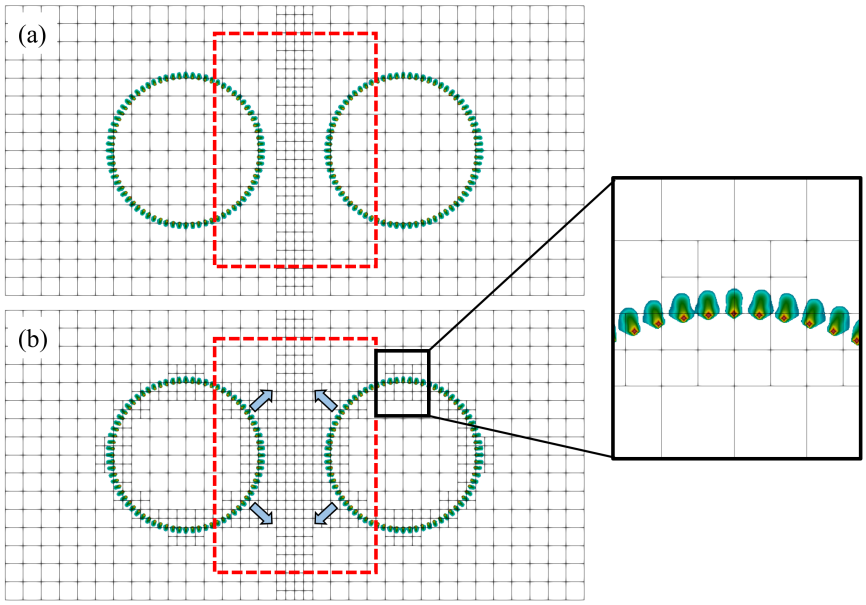


Figure 6.9.: Variable mesh grids used to refine a) the reconnection layer only, and b) the arrays as well as the inflows entering the layer. In the latter, the $\mathbf{j} \times \mathbf{B}$ force vectors are shown to demonstrate the region of influence that is refined besides the array and layer. Inset shows the wire cores within the high resolution region, while ablated plasma enters low resolution areas. Red dashed boxes indicate regions of interest for further discussion.

assess the impact of the layer and array resolution separately.

The following discussion compares results from a uniform grid mesh at 50 μm, a mesh that resolves only the layer to 25 μm (Fig 6.9(a), and a mesh that additionally resolves the arrays and ablation streams at 25 μm (Fig 6.9(b)). As a point of reference, the uniform mesh at $\Delta x = 50$ μm uses 512 processors, the layer-resolved mesh uses 608, the layer-and-array-resolved mesh uses 872, and a uniform mesh at $\Delta x = 25\mu$m uses 2048.

Uniform mesh vs. resolved layer

Resolving only the central layer shows some morphological changes, however key parameters such as the number of plasmoids formed or collapsed layer temperature are mostly similar. The left two images in Fig 6.10 shows the layer plasma which reaches $100 \leq T_i \leq 2000$ eV, for both simulations. The mesh is overlaid in grey for the variable resolution case. Lineouts through the central columns (shown as a red line on the left plots) are presented on the right, for three different times. The left images and the topmost right panel both correspond to 100 ns, the central and bottom righthand panels are 50 ns and 100 ns later respectively.

The left images show peaks in temperature at $T_i \sim 1$ keV where interaction regions between ablation streams have reflected of the axis and are thermalised. The top and bottom ends of the layer at $T_i \sim 2 - 2.5$ keV correspond to low-density outflows accelerated to >200 km/s. The inflow conditions are nearly identical except for small

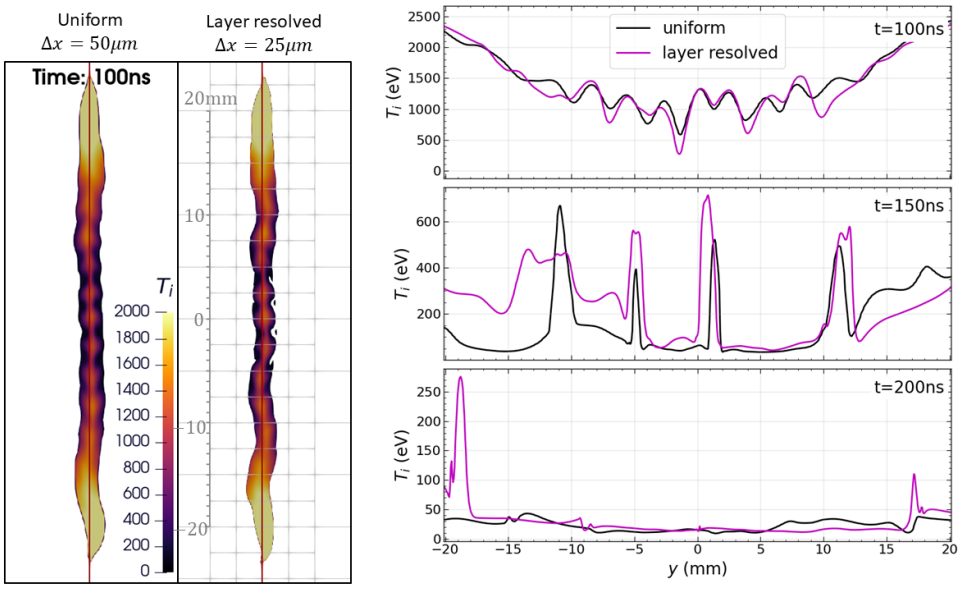


Figure 6.10.: Effects of resolving only the reconnection layer at 25 μm, at 100 ns. *Left:* images showing the layer-only region showing $T_i > 200$ eV, where image on the right has processor splitting overlaid in grey lines. *Right:* lineouts through the $\hat{y}$-axis for both simulations showing ion temperature at three different times: (top to bottom) 100, 150, and 200 ns.

variations from the temperature variations, and consequently the peaks in ion temperature reach similar magnitudes. Despite similar inflows, the higher resolution layer is more susceptible to left-right variations therefore exhibits more shear in the reflections. As a result, the lineouts through the centre do not coincide with the peak values, and a comparison artificially shows more difference. 50 ns later both have formed plasmoids as shown in Figs 6.7 and 6.8 at 150 ns. The peak ion and electron temperatures are 50-100 eV higher in the high resolution simulation, which increases the thermal pressure in the plasmoids and resists compression from the inflow for marginally longer. Therefore the rise in plasmoid density lags behind the low resolution case by 10-20 ns. This is reflected in the time of plasmoids being extinguished, which occurs at 180 ns for the low resolution and 200 ns for the high resolution simulation, although this difference may be an overestimate considering the temporal resolution. Finally, at 250 ns when both simulations are in their cooled steady state, the average temperature across both layers are of the order of 25 eV.

Resolved layer vs. resolved layer and arrays

The behaviour when also resolving the wires at 25 μm is described here, i.e. between the simulation splittings shown in Figs 6.9(a) and (b). These are referred to as the layer-only vs. layer-array case to describe which regions have been resolved at the higher resolution. The comparison between these is less straightforward than above, as the wire initialisation is different in the layer-array case. Here, wires are initialised as 4×4 cell cores in the layer-array (rather than 3×3 as in the layer-only case)

giving a density scaling of $f = \pi r_w^2/16\Delta x\Delta y$ and therefore $\rho_0 = f\rho_s = 2120$ kg/m^3, significantly closer to ρ_s than in the layer-only case. However, this has implications for other wire properties. Specifically, the use of LMD tables leads to a density dependence of conductivity. Therefore the layer-array case is initialised at higher (and closer to solid) conductivity than in the layer-only case. The effect of switching transport models was shown in Fig 4.10 in the context of 1D single wire behaviour (§4.3). While not an identical comparison, the Spitzer model simulation was similarly initialised at lower conductivity than the reference LMD case, and showed an earlier expansion due to stronger Ohmic heating of the surface as a result of a longer diffusion time through the wires. This effect, combined with the more diffusive expansion of material at low density, leads to faster rate of expansion from the wire cores compared to the layer-only case. The density profiles at 50 ns of a segment of arrays from both simulations are shown in Fig 6.11, as well as the respective initialisation of wire cores.

While this behaviour results in differences between the timings of the arrays, the characteristics of the streams emanating from the wires cores are similar. A comparison of electron temperature profiles at 150 ns is shown in Fig 6.12 in terms of inflow, and the central current sheet within the layer. Also shown on the right of the two plots are lineouts spanning 16 mm in the $\hat{x}$-direction across the brightest plasmoid for each case. The streams are 7 eV in layer-only case, and 8 eV in the layer-array case, however

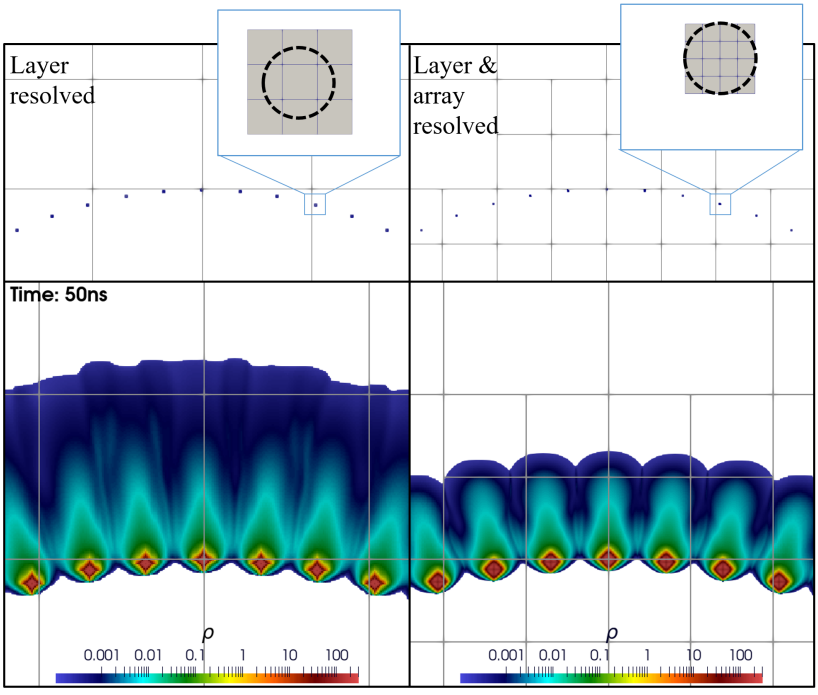


Figure 6.11.: Effect of ablation rate using wires initialised at 50μm resolution using 3×3 cells (left), and initialised at 25μm using 4×4 cells (right). Top images show segment of arrays at initialisation, where the wire spacing is not seen to be strongly affected. Bottom images compares density (in kg/m^3 of ablation streams at 50 ns for both resolutions. Grid splitting is shown overlaid in grey lines.

the shock interactions reach higher densities and lower temperatures in the latter as a result of lower numerical diffusion. This enhanced interaction leads to larger deflections in velocity in the $\hat{y}$ direction. As a result, a larger number of 'unstable' regions within the layer undergo plasmoid formation. In the lower resolution and more radially directed case, the plasmoids that do form reach higher temperatures. This behaviour is reflected in the plots on the right of Fig 6.12, where the peak values of T_e in the layer-only and layer-array cases are 358 eV and 435 eV respectively. Finally, the electron temperature within the plasmoids drop to below 20 eV (from > 350 eV) at 20 ns after the equivalent layer-only case. This delay in part due to the temporal offset from the beginning, but also due to a persistently higher cooling rate from the layer in the high density layer-only case.

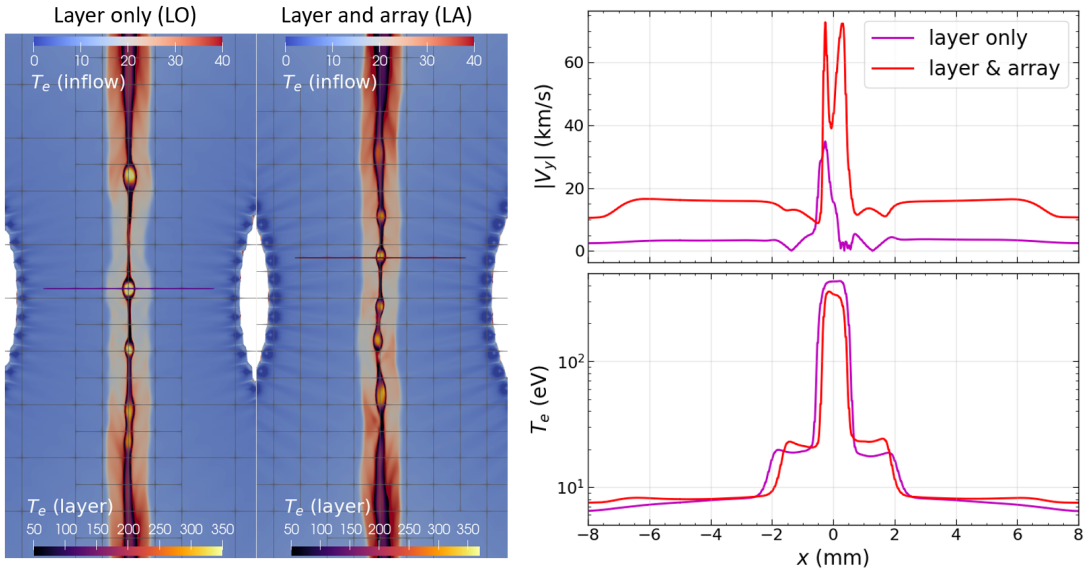


Figure 6.12.: Effect of inflow and layer conditions comparing wires initialised at 50μm versus 25μm resolution. *Left:* Images of electron temperature in eV at 150 ns. Two sets of colour scales used for clarity; blue-red (0-40 eV) for inflow, and purple-yellow (50-370 eV) for central current sheet within layer. Lineouts spanning 16mm across high temperature plasmoids for 50μm (magenta line) and 25μm (red line) resolution shown on the right, in terms of y-component velocity in $\hat{y}$-direction, and electron temperature.

6.3.4. Radiation effects

This section assess the difference between different radiation models on the reconnection dynamics. Unlike using variable resolution in different regions of the grid, the radiation models cannot be applied on the layer and wires separately and their effects cannot be as readily isolated. However, as expected from the effects of varying radiation model on the single wire system (§4.3), the wire ablation is highly sensitive to this model option and therefore changes the inflow into the layer. Note, simulations presented in this

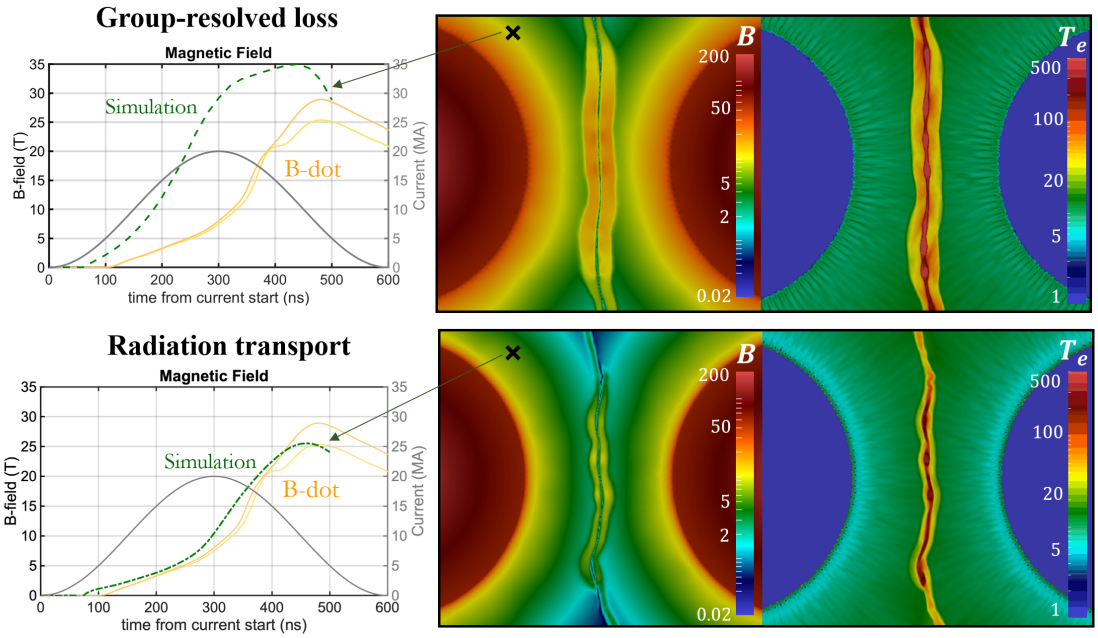


Figure 6.13.: 2D simulations of reconnection dynamics comparing group-resolved radiation loss model and radiation transport model. *Left:* plots comparing experimental (yellow lines) and synthetic (green lines) B-dot probes from model options, placed 5mm away from array. Current traces (grey lines) shown on right axes. *Right:* profiles of magnetic field in T (left images) and electron temperature profiles in eV (right images) from each simulation, at 200 ns. Simulations run by Prof. J. Chittenden, experimental B-dot data courtesy of Sandia National Laboratories.

section [2] were run using 150 wires and using TF-EoS and Spitzer transport instead, leading to differences in plasma properties but overall exhibiting the same trends.

Fig 6.13 compares two simulations run with each model option at 50 μm resolution. The top row shows results from using the group-resolved model (as was used in previous results), and the bottom row showing radiation transport results. The leftmost images show profiles of magnetic field versus time using experimental (two yellow lines) and synthetic (green line) B-dot probes, overlaid with current traces (grey line). The profiles from the radiation transport case are observed to give significantly better agreement with experiment data in two respects; firstly the loss case undergoes a relatively sharp increase, and secondly therefore overestimates the peak amplitudes by > 5 T even accounting for the temporal offset. The radiation transport case on the other hand undergoes a gradual increase, with an inflection at ~ 300 ns and eventually peaks at 25T. The experimental B-dot shows an inflection 50 ns later at 350 ns, however this discrepancy is equal to the timing difference in the start of detecting magnetic field, suggesting it could be a systematic offset.

Images in the centre and right-hand columns of Fig 6.13 show the simulation profiles of the magnetic field and electron temperature using both models, at the same time (200

[2] Simulations run by Prof. J. Chittenden, analysis performed by the author.

ns). The approximate location of the synthetic B-dot (obtained by placing a probe in a specified location and extracting data), is indicated by a black cross. The magnetic field profiles in the radiation loss case show significantly higher magnitudes, both around the wire and piled up at the layer. The layer width is also significantly larger, whereas the in the radiation transport case it is thinner and more susceptible to perturbations across the inflow. The difference in value around the array can be explained by considering the electron temperature. At early times, the ablation streams from both cases heat to approximately similar temperatures, 10-15 eV. By 200 ns, the density close to the wires cores are increased in the radiation transport case, leading to a higher rate of loss. The streams cool to $\sim$ 5 eV, as seen in the electron temperature profile in Fig 6.13. The magnetic Reynolds number of the flow ($R_M \propto 1/\eta$) drops from the increased resistivity, and less field is transported with the inflow. The increased density of ablation streams around the wire core results in a transition to higher optical depth at 250 ns, allowing the radiation to be trapped and re-heat the flow. This is reflected in the profile of R_M versus time (plotted in Fig 14 in Ref [77]), which falls and rises with electron temperature, and equivalently in the inflection of the B-dot profile at 300 ns (Fig 6.13).

Focusing on the layer itself, a region of localised heating around the layer is apparent

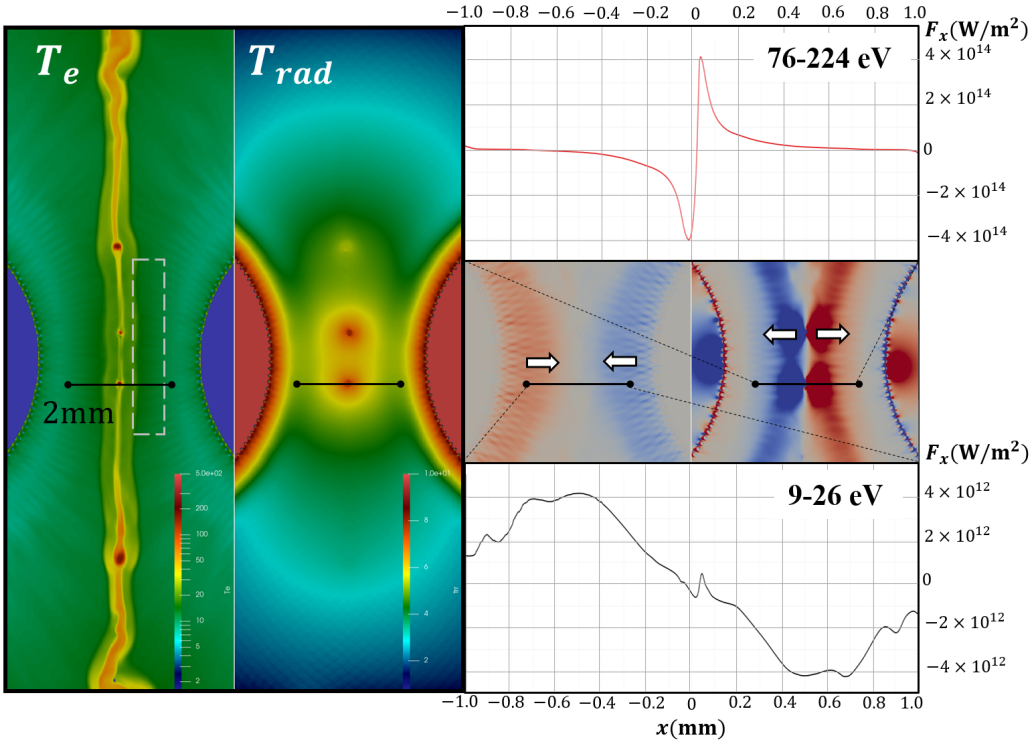


Figure 6.14.: Effect of using radiation transport, shown at 240 ns. *Left:* profiles of electron and radiation temperature in eV. The radiation temperature shows a diffusive wave from the plasmoids, leading to a preheat of the inflow, shown for the right array in grey dashed box). *Right:* Radiation flux in $\hat{x}$-direction for two energy groups, 9-26 eV (left central image and bottom lineout) and 76-224 eV (right image, top lineout).

at the time of peak plasmoid production, and is only observed in the radiation transport case. This is highlighted for the right-hand wire array in the dashed grey box around the leftmost image in Fig 6.14, showing the electron temperature profile 40 ns after the state in Fig 6.13. Shown to the right is the radiation temperature profile, which peaks at 8.5 eV at the centre of the plasmoid and decays to 4-6 eV in the surrounding regions. The 5.5 eV contour that surrounds the central two plasmoids is particularly visible, and corresponds roughly to the area of increased localised heating in the ablation streams. The non-local re-absorption of radiation from the plasmoids therefore introduces a radiative precursor in the inflows. The images on the right show the behaviour of the radiative fluxes in the $\hat{x}$-direction, for two out of ten spectral groups used (9-26 eV and 76-224 eV). The central images show their profile on the same spatial scale as the temperature images, where positive values are shown in red and negative values in blue, while the plots above and below show lineouts spanning 2mm across the lower central plasmoid. The lower energy range (left image, bottom plot) shows radiation transferred from the wire cores radially outwards, as indicated by the direction of the white arrows. In the higher energy range (right image, top plot) radiation is transported outwards from the plasmoids with larger fluxes by two order of magnitudes, providing the non-local heating of the ablation streams as they enter the reconnection layer. This effect is clearly not captured using the local, group-resolved loss model.

6.4. 3D Simulations

To conclude this chapter, a preliminary 3D simulation of wire array-driven reconnection is presented with the layer resolved at 25 μm (run using the group-resolved loss model). Simulations have previously been performed at 50 μm resolution, this the first at this resolution. While the behaviour of the reconnection layer is expected to be fairly uniform

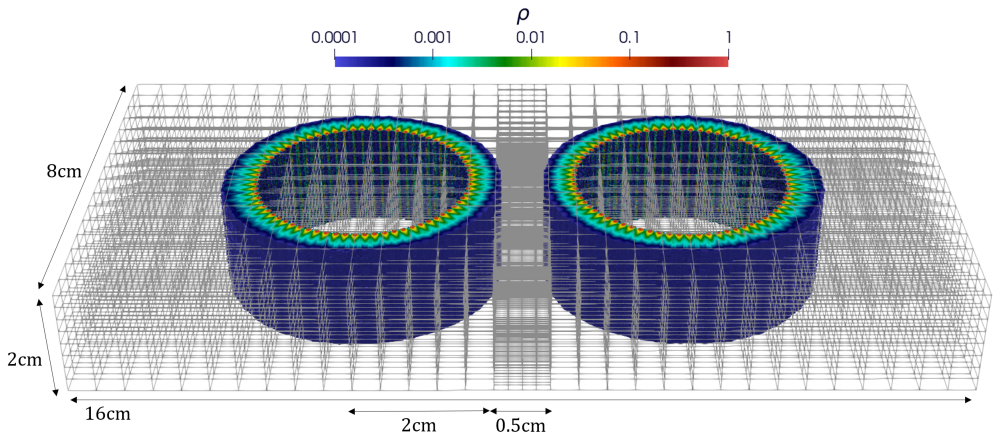


Figure 6.15.: Mesh splitting for 3D simulation of wire array reconnection, showing density profile (kg/m^3) from wires at 50 ns, and processor splitting in grey outlines. Dimensions of full domain, array radius, and high resolution width provided.

in the axial extent and therefore quasi-2D, single wire simulations demonstrated how seeding perturbations at even cell-size resolution can lead to the non-linear development of large scale instabilities, which has consequences for plasmoid formation.

A brief review is provided here of the previous 3D simulations, described in Ref [77]. Firstly, these revealed quasi-2D behaviour as expected, with some variation in the $\hat{z}$-direction. Notably, the plasmoids formed 'flux-ropes' which can be considered the 3D analogue of the 2D feature [77]. These are isolated columns of high current, that are susceptible to $m = 1$ instabilities. The 3D simulations were compared against 2D simulations using different radiation models, where the radiation-driven dynamics was deduced to be the biggest lever in reducing the layer width.

The simulation presented here serves two purposes, firstly it verifies the behaviour observed at a higher resolution, and secondly it is used to compare like-for-like simulations performed in 2D and 3D. The set-up described in §6.3.1 for the layer-only resolved case is extended here to an axial extent of $L_z = 2$ mm, shown in Fig 6.15. To expedite simulation run time, the number of cells per processor in the axial direction was reduced relative to the other two directions, utilising 5888 processors in total. It was run until 90 ns before terminating, due to stricter timestep restrictions at the higher resolution, particular when the layer outflows reach several keV ion temperatures. However, some key observations can be made up until this time.

The first is the profile of the layer in the axial direction. At this time, the inflows have been in contact with the central axis for 30 ns, i.e. not enough time for the layer to experience strong mass and magnetic field pile up, but enough time for the radially diverging inflow to form a layer that is 4mm long, or $L_y/2$. Therefore the effect of axial modulations from the wires will be evident. The state is shown in Fig 6.16. The top image shows the axial extent of the wire cores in white, with slices of density shown in the $\hat{x} - \hat{y}$ plane and $\hat{y} - \hat{z}$ planes. The bottom two images show the current density also along the $\hat{y} - \hat{z}$ plane, in terms of its magnitude and $\hat{x}$ (out of plane) component. Overall, the same features as described in Ref [77] are observed, namely dense, current-carrying flux ropes that are nearly uniform. The modulations along the axial direction are visible in the current density, where wavelengths of the order of $\lambda \approx L_z/10 = 200$ μm are observed. While these modulations appear to break up the flux ropes, this is an artifact of taking a single slice, rather than tracing the current density path. The $\hat{x}$ component shows the current weaving in and out of the page in a near-helical manner suggestive of the $m = 1$ instability (also observed at 50 μm [77]). These distinct structures are elongated in the $\hat{y}$-direction as they get compressed by the ram pressure of the inflows and advected with the outflows.

The second observation is made when comparing 2D and 3D simulations from a top-down view, shown in Fig 6.17 in terms of the electron temperature. While the flow profiles are nearly identical, as expected without changing any other parameters, the reconnection layer is more uniform in the 3D rather than 2D case, while reaching visibly lower peak temperatures.

This can be understood as follows. The reduced dimensionality case inherently as-

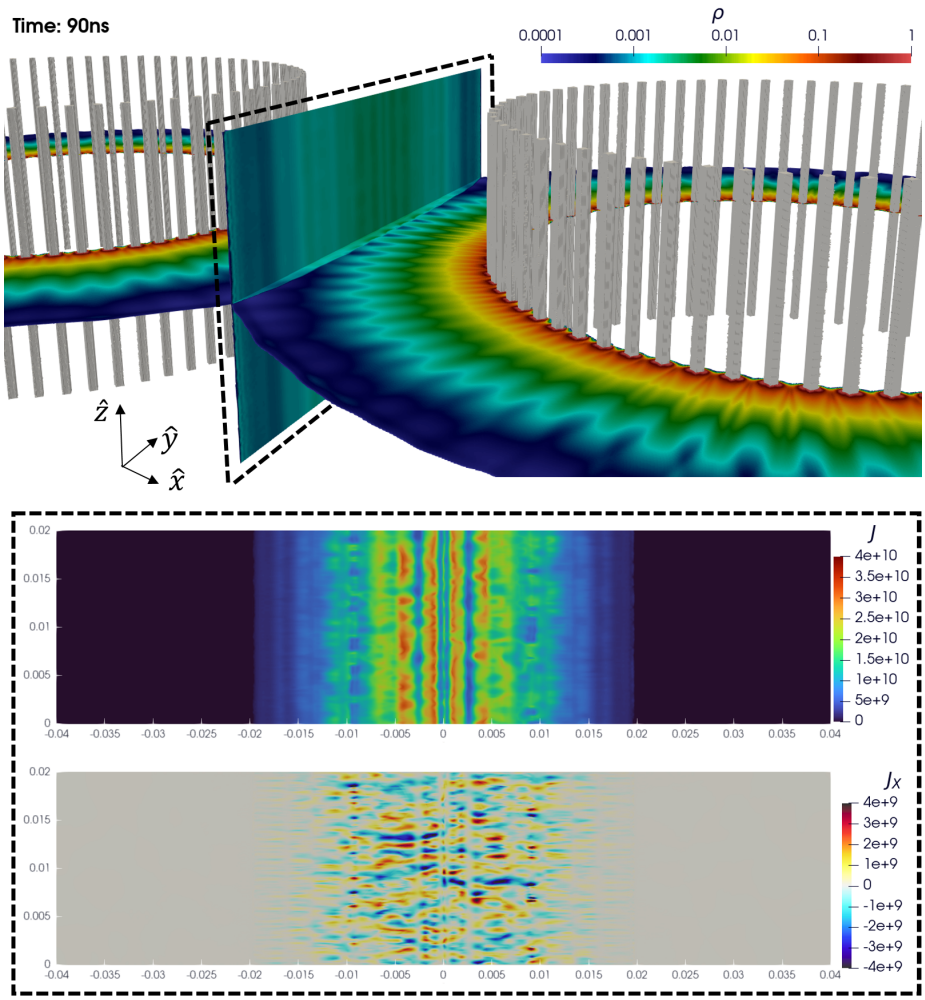


Figure 6.16.: *Top:* 3D view of ablation streams from wire arrays, showing wire cores in grey, and slices through the density profile (in kg/m^3) in the $\hat{x} - \hat{y}$ and $\hat{x} - \hat{z}$ directions. *Bottom:* slice through $\hat{x} - \hat{z}$ plane showing magnitude (top) and $\hat{x}$-component of current density (in A/m^2).

sumes that any perturbations in the $\hat{x} - \hat{y}$-plane extend infinitely in the axial direction, therefore are effectively amplified compared to the 3D case. In 3D, if a single 'layer' of ablated material expands slower faster than the layer above or below it, the pressure gradient acts to expand flow in the axial direction, thereby smoothing the leading edge. The v_z component is smaller than v_x by a factor of 10-100 during the ablation phase, however further isotropisation occurs once the flows have collided and stagnated on axis. The higher peak temperatures in 2D are a result of larger radial component (v_x or v_y) for a given velocity vector in the absence of an additional degree of freedom.

6.5. Concluding remarks

The simulation studies performed in this section have demonstrated the applicability of mesh refinement to magnetic reconnection achieved on a pulsed-power driven platform. Simulations performed in 2D demonstrate that while resolving only the reconnection

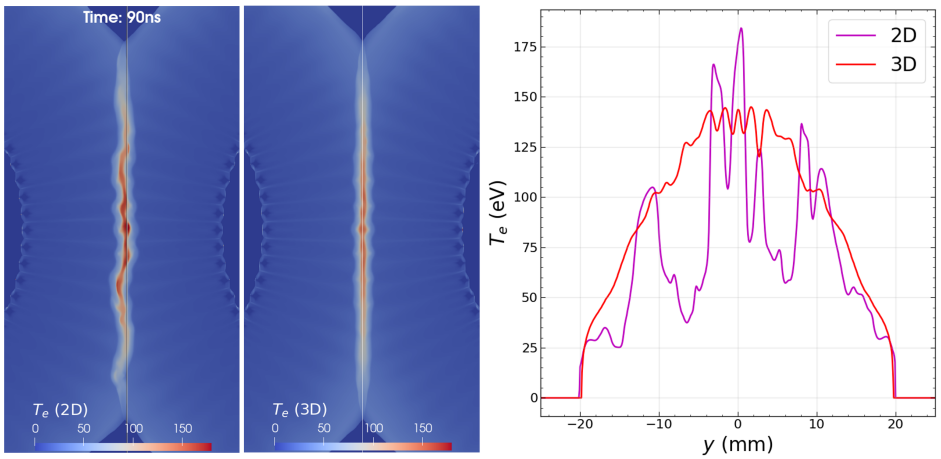


Figure 6.17.: Top-down view of electron temperature profiles (in eV) of reconnection layer for a 2D (left image) and 3D (right image) simulation using equivalent variable resolution configuration. The plot on the right shows lineouts through $\hat{y}$-axis for both images.

layer at higher resolution can affect its uniformity, resolving the wires and ablation streams into the layer make a more substantial difference. Specifically, the irregularities due to shock interactions between ablations streams from individual wires are enhanced at higher resolutions, leading to variations in the velocity vectors that lowers the shock heating of the colliding flows, and therefore produces lower temperature plasmoids. However, both cases of resolving the wires at 25 μm and 50 μm requires a scaling of the initial wire density to match the linear mass at solid density. The simulations were performed using an LMD transport model, therefore the conductivity has a density dependence, and therefore results from the two resolutions are not directly comparable. However, this does indicate that the system is highly sensitive to wire array dynamics. Use of radiation transport instead of the group-resolved loss model additionally changes the dynamics, in both the reconnection layer and the wire array. The radiation transport simulations show initially cooler plasma surrounding the wire cores that carry lower magnetic field with the ablation streams, but increase in density later in time, trapping the radiation and increasing the magnitude of field advected. Finally, 3D simulations reveal a smoothing of the layer in the $\hat{x} - \hat{y}$ due to perturbations in the axial direction, which additionally reduces the plasmoid temperature.

The above investigation reveals a complex system where several contributing factors can affect the inflow conditions into the layer, which then influences the transition from plasmoid-driven to radiatively-dominated reconnection. The coupling of different effects, for example resolving the wire cores at high resolution while also using radiation transport, is expected to lead to yet different dynamics as presented above, and will require the mesh refinement capability to access these resolutions. Future work will look to validate the ablation stream behaviour using higher fidelity models from cold-start conditions against experimental data.

7. Conclusions

The work presented in this thesis has demonstrated the use of a static mesh refinement capability into the radiation-magnetohydrodynamics (rad-MHD) code Chimera to study pulsed-power driven HEDP systems that undergo strong radiative cooling. Previous modelling efforts to study plasmas undergoing radiative collapse have been resolution limited; these simulations are now accessible using mesh refinement, for a fraction of the computational cost as using an equivalent uniform mesh. A review of the capability is first provided, followed by discussion of the simulations studies performed using it.

7.1. Mesh refinement capability

The primary outcome of this thesis was the development of a static mesh refinement (SMR) capability using block-based refinement for use in an MHD code with a constrained-transport solver. There exist several refinement strategies for delivering targeted resolution to regions of interest within a simulation domain. A simple yet robust approach strategy was adopted, based on an earlier implementation by previous developers of the code, but extended to an arbitrary number of refinement levels and with flexible directional splitting. In this method, all processors are generated with the same number of computational cells in each spatial dimension (and therefore total number of cells), however the resolution of cells in each processor can vary. This has two key advantages: firstly, load balancing is achieved by default, and secondly, no modifications are required for the physics solvers themselves. However, three other additional functionalities are required: 1) the procedure for constructing grid topology, 2) an output format based on the grid topology, and 3) a procedure to transfer ghost cell variables across processors of different resolutions. The capability developed provides standalone modules for each functionality, which can be imported into other resistive MHD codes. Additionally, a novel algorithm was developed to perform ghost cell transfers for MHD variables for a resistive, constrained-transport (CT) solver that is divergence-free and curl-preserving, and utilises non-blocking MPI communications to improve performance.

Future work

The capability was demonstrated on various large-scale, integrated simulations. However, there remain key avenues for improvement. Firstly, grid topology is currently constructed using simple geometries to refine regions of interest, such as cuboids and cylindrical shells. These are appropriate for a large range of pulsed-power systems such as single wires and wire arrays respectively. However, incorporating an STL reader to

define arbitrary geometries would enable splitting around more complicated structures, for example the helical return posts used in the 'dynamic screw pinch' (DSP) configuration of MagLIF. Secondly, the grid output format current uses a file-per-processor approach, which can produce tens of thousands of files for a large, 3D simulation. An alternative solution is using a VTK unstructured grid format, which stores all variable data in one file, although this also requires storing the grid topology, which leads to considerable redundancy for a static topology. A further extension to this capability is to develop a suite of post-processing tools e.g. to generate synthetic diagnostics. This is coupled with the choice of grid output format. For example, a commonly used procedure is integration through a specified line-of-sight, as was performed for synthetic self-emission images of 2D solid X-pinch simulations. This was achieved by interpolating the data on a high resolution uniform grid, however this is unfeasible for large 3D data. A future capability may use a single-file, unstructured grid format, however for large simulations this would require loading the data incrementally to perform the integration, as storing the data at once would exceed the memory limitations of available computing facilities.

Finally, a possible extension of this implementation from a static (grid topology generated at initialisation) to an adaptive (grid topology evolves with grid variables) capability is discussed. An increasing number of MHD codes use adaptive mesh refinement (AMR) libraries to provide dynamic grid splitting. Extending the existing implementation to an adaptive one requires considerations of all three aforementioned functionalities: the grid splitting criteria would be redefined in terms of grid data variables, the output format would require storing the grid topology every time it was modified, and even the prolongation/restriction operators would require bespoke treatment. Finally, load balancing would need to be directly managed, for example using knapsack algorithm on variable block sizes as performed in existing libraries. Various libraries exist that are optimised to perform the above, integrating one (such as AMReX) in Chimera is likely to be more appropriate than re-purposing the current capability.

7.2. Radiatively cooled systems

The behaviour of collapsing of Z-pinches down to small radii were studied using single wires and X-pinches to create the final state. Three sets of experiments were simulated: two on $\sim$200 kA generators, and one on a 6 MA generator. Chapter 4 presented simulations of aluminium single wire explosions on a 150 kA generator. Despite the simplicity of the system, these simulations exhibited complicated dynamics with strong sensitivity to various model options. Particularly, use of a higher fidelity transport model showed better matched core expansion rates to experimentally observed rates from Schlieren imaging. The compression of the corona at later times was several times smaller using loss models versus radiation transport, due to radiative preheat raising the core thermal pressure. However, the latter is substantially more computationally intensive, and therefore intractable to run on the larger 2D or 3D simulations, also presented. These

showed the formation of hotspots at the necks of $m = 0$ instabilities, which were disrupted before reaching high pressures. Simulations performed in 3D compared to 2D resulted in a drop in peak pressure by an order of magnitude.

Chapter 5 presented results from simulations of two different X-pinch experiments: 'typical' wire X-pinches (WXP) performed a lower current generator, and the solid X-pinch (SXP) configuration performed on the 6 MA SATURN generator. The geometry of X-pinches focuses the current to central crossing point, and therefore significantly higher hotspot pressures than in single wire simulations were reached in these systems, even at comparable currents. Previously unobserved 3D behaviour was shown in the WXP simulations, which produced an asymmetric morphology of the crossing point at early times due to the offset of the wires from the centre, but that later formed a uniform Z-pinch. The SXP simulations were performed in 2D due to their inherent azimuthal symmetry, however the conditions in these experiments were considerably more extreme and therefore also challenging to simulate. The higher currents produced hotspot pressures of up to several hundred terapascals, comparable to the conditions in capsule implosions on laser facilities. Again, significant sensitivity was observed to the compressed state using different radiation models.

The hotspot conditions produced in all three systems were extreme, and some similarities were observed. Most notably, these all showed a high pressure, relatively cold, central region, surrounded by a current-carrying edge plasma that had higher temperatures due to strong Ohmic heating and relatively lower densities (although still distinct from the hot, tenuous coronal plasma). The radiative losses from edge regions in all simulations were consistently observed to be higher than from the core. Additionally, the brightest emission regions from synthetic self-emission images of the SXP simulations had radii corresponding to the edge plasma boundary, as radiation from the central core had too high an optical depth. This suggests that measurements of plasma parameters from bright hotspots along several Z-pinch systems (including X-pinches) may not be from uniform conditions.

Finally, simulations performed of the magnetic reconnection experiments in Chapter 6 demonstrated the benefit of mesh refinement when applied to large-scale systems that form smaller scale interactions regions. Simulations of the MARZ experiments required resolving two wire arrays spanning several centimeters, and simultaneously ~ 50 μm scale reconnection dynamics in the central layer, which underwent radiative cooling to small widths. Numerical investigations of applying mesh refinement to different regions of this system in 2D show that modulations in the wire array are enhanced by resolving the wire cores to 25 μm, which has potential implications for plasmoid (flux-rope) formation in 3D. The advection of magnetic field into the layer was observed to be lower when using radiation transport versus using loss models, and underwent a transition as plasma around the wire cores become increasingly optically thick. The interplay of high resolution and radiation transport are also expected to effect dynamics, and will be explored in future studies.

Future work

The results presented in this thesis demonstrate how a mesh refinement capability can be used to study previously inaccessible regimes. However, this tool can be used in numerous applications, particularly for ICF and MIF applications where systems undergo strong convergence. For example, preliminary simulations are being performed in 2D cylindrical geometry of MagLIF liners, where high resolution is applied on the liner surface to study MRT growth, and also on the axis to observe the impact on yield. Extended simulations in 3D could also be performed of the MagLIF system include the DSP return posts, if an STL reader is implemented as described above.

With respect to the systems modelled in this work, there is a clear need to progress the simulations described above further past the time of peak compression. In the wire X-pinch case this was primarily hindered by low density plasma reaching the domain boundary, which is simply resolved using large domain boundaries as enabled using mesh refinement. In the SXP case, the high Alfvén speeds that form in low-density plasmas, particularly close to peak current, resulting in prohibitive timesteps that are further restricted by a large number of required thermal sub-cycles to resolve strong Ohmic heating in these regions. Reducing the density floor to overcome this is unlikely to alter the dynamics in the SXP case as the current path preferentially followed the solid density, high conductivity regimes (as produced using LMD tables). However, the coronal plasma is important in thin wire systems, requiring an alternative solution. It has been suggested that extended-MHD effects such as the Nernst and Ettingshausen terms may reduce Ohmic heating in low density regions; these will be explored in future.

Finally, the broader dynamics of the above systems were explored, however the simulation set-up could be used to run parameter studies. For example, at the time of writing, there is flexibility in the design of remaining experiments in the MARZ campaign. Simulation studies are underway to assess the effect of wire array separation on the layer temperature, which further extends the simulation domain and therefore further necessitates the use of mesh refinement. The wire X-pinch has been extensively characterised experimentally, similar studies could be performed numerically and compared against existing data.

The discussion above has demonstrated the applicability of a robust mesh refinement capability to several pulsed-power system of interest, and shown the potential to perform numerical studies at previously inaccessible resolutions.

A. Algorithms for MHD solvers

The algorithm below describes the operator-split vector potential solver in Chimera, which solves for $\partial \mathbf{A}/\partial t$, then updates the magnetic field using $\mathbf{B} = \nabla \times \mathbf{A}$.

The advection step is first performed on the hydrodynamic timestep Δt, followed by resistive subcycling performed on the MHD timestep, Δt_{MHD}, which resolves dynamics at the speed of light c (typically reduced to $\sim c/100$ for expediency).

Algorithm 1 Vector-potential solver

! - *advection step*
if advection **then**

(apply **B** boundaries)
Update **A** advective - Δt
$> \mathbf{A}^n = \mathbf{A}^{n'} + \Delta t \left[\mathbf{v} \times \mathbf{B}^{n'}\right]$

Update **B** from **A**
$> \mathbf{B}^n = \nabla \times \mathbf{A}^n$

end if

! - *resistive sub-cycling*
if resistive **then**
while $i_{\text{sub}} <= n_{\text{subcycles}}$ **do**

(apply **B** boundaries)
Update **A** resistive from $\nabla \times \mathbf{B}$ and **H** - $\Delta t_{\text{MHD}}/2$
$> \mathbf{A}^{n+1} = \mathbf{A}^n - \Delta t_{MHD} \left[\mathbf{E}^n e^{\kappa/2} - \nabla \times \mathbf{B}^{\mathbf{n}}(1 - e^{\kappa/2})\right]$

Update **B** from **A**
$> \mathbf{B}^{n+1} = \nabla \times \mathbf{A}^{n+1}$

(apply **B** boundaries)
Update **E** diffusive - Δt_{MHD}
$> \mathbf{E}^{n+1} = \mathbf{E}^n e^{\kappa} + \nabla \times \mathbf{B}^{n+1}(1 - e^{\kappa})$

end while
end if

! - *update current*
$\mathbf{j}^{n+1} = \mathbf{E}^{n+1}/\eta$

The algorithm below describes the constrained-transport solver in Chimera, which solves for $\partial\mathbf{E}/\partial t$, then updates the magnetic field using $\partial\mathbf{B}/\partial t = -\nabla \times \mathbf{E}$. The term $\mathbf{E}_{\rm temp}$ is used as a temporary variable to store magnetic field 'fluxes' (equivalent to $\mathbf{A}$), while the electric field is evolved separately. In both algorithms, $\kappa = -\mu_0 c^2 \Delta t_{\rm MHD}/\eta$.

Algorithm 2 New constrained-transport solver

! - *advection step*
if advection **then**

(apply **B** boundaries)
Update $\mathbf{E}_{\rm temp}$ (advective)
$> \mathbf{E}_{\rm temp} = \mathbf{v} \times \mathbf{B}^{n'}$

(apply $\mathbf{E}_{\rm temp}$ boundaries)
Update **B** from $\mathbf{E}_{\rm temp}$ - $\Delta t_{\rm MHD}$
$> \mathbf{B}^{n} = \mathbf{B}^{n'} + \Delta t_{\rm MHD}\,[\nabla \times \mathbf{E}_{\rm temp}]$

end if

! - *resistive sub-cycling*
if resistive **then**
while $i_{\rm sub} <= n_{\rm subcycles}$ **do**

(apply **B** boundaries)
Update $\mathbf{E}_{\rm temp}$ (resistive) - $\Delta t_{\rm MHD}/2$
$> \mathbf{E}_{\rm temp} = \mathbf{E}e^{\kappa/2} + \eta\nabla \times \mathbf{B}^{n}(1 - e^{\kappa/2})$

(apply $\mathbf{E}_{\rm temp}$ boundaries)
Update **B** from $\mathbf{E}_{\rm temp}$ - $\Delta t_{\rm MHD}$
$> \mathbf{B}^{n+1} = \mathbf{B}^{n} + \Delta t_{\rm MHD}\nabla \times \mathbf{E}_{\rm temp}$

(apply **B** boundaries)
Update **E** diffusive - $\Delta t_{\rm MHD}$
$> \mathbf{E}^{n+1} = \mathbf{E}^{n}e^{\kappa} + \eta\nabla \times \mathbf{B}^{n+1}(1 - e^{\kappa})$

end while
else
$\mathbf{E}^{n+1} = \mathbf{E}_{\rm temp}$ (advective)
end if

! - *update current*
Update $\mathbf{j}^{n+1} = \mathbf{E}^{n+1}/\eta$

Bibliography

[1] N. Bennett, D. R. Welch, C. A. Jennings, E. Yu, M. H. Hess, B. T. Hutsel, G. Laity, J. K. Moore, D. V. Rose, K. Peterson, and M. E. Cuneo. Current transport and loss mechanisms in the z accelerator. *Physical Review Accelerators and Beams*, 22:120401, 12 2019. 8, 20, 21

[2] R.J. Leeper and et al. Z-pinch driven inertial confinement fusion target physics research at sandia national laboratories. *Nuclear Fusion*, 39:1283, 9 1999. 8, 22

[3] M. R. Gomez, S. A. Slutz, A. B. Sefkow, D. B. Sinars, K. D. Hahn, S. B. Hansen, E. C. Harding, P. F. Knapp, P. F. Schmit, C. A. Jennings, T. J. Awe, M. Geissel, D. C. Rovang, G. A. Chandler, G. W. Cooper, M. E. Cuneo, A. J. Harvey-Thompson, M. C. Herrmann, M. H. Hess, O. Johns, D. C. Lamppa, M. R. Martin, R. D. McBride, K. J. Peterson, J. L. Porter, G. K. Robertson, G. A. Rochau, C. L. Ruiz, M. E. Savage, I. C. Smith, W. A. Stygar, and R. A. Vesey. Experimental demonstration of fusion-relevant conditions in magnetized liner inertial fusion. *Physical Review Letters*, 113:155003, 10 2014. 8, 22

[4] M. G. Haines, S. V. Lebedev, J. P. Chittenden, F. N. Beg, S. N. Bland, and A. E. Dangor. The past, present, and future of z pinches. *Physics of Plasmas*, 7:1672–1680, 5 2000. 8, 28

[5] Jon Bashor. Berkeley lab to lead amr co-design center for doe's exascale computing project, 11 2016. 8, 29

[6] Jon Larsen. Foundations of high-energy-density physics: Physical processes of matter at extreme conditions. *Foundations of High-Energy-Density Physics: Physical Processes of Matter at Extreme Conditions*, pages 1–743, 3 2017. 8, 40

[7] Steffen Faik, Anna Tauschwitz, and Igor Iosilevskiy. The equation of state package feos for high energy density matter. *Computer Physics Communications*, 227:117–125, 6 2018. 8, 43, 44

[8] Y T Lee and R M More. An electron conductivity model for dense plasmas. *Citation: The Physics of Fluids*, 27:1273, 1984. 9, 45, 47

[9] Howard Scott. Lecture 17 – radiation transport. *High Energy Density Spectroscopy: Ions, Electrons & Photons (lecture series)*, 2023. 9, 50, 51

[10] J.D. Pecover. *Instability Growth for Magnetised Liner Inertial Fusion seeded by Electro-thermal, Electro-choric and Material Strength Effect*. PhD thesis, Imperial College London, 2015. 10, 66, 67

[11] Dinshaw S. Balsara. Divergence-free adaptive mesh refinement for magnetohydrodynamics. *Journal of Computational Physics*, 174(2):614 – 648, 2001. 11, 80, 81

[12] D. A. Hammer and D. B. Sinars. Single-wire explosion experiments relevant to the initial stages of wire array z pinches. *Laser and Particle Beams*, 19:377–391, 2001. 12, 92, 95

[13] Edmund P. Yu, B. V. Oliver, D. B. Sinars, T. A. Mehlhorn, M. E. Cuneo, P. V. Sasorov, M. G. Haines, and S. V. Lebedev. Steady-state radiation ablation in the wire-array z pinch. *Physics of Plasmas*, 14:022705, 2 2007. 12, 92, 93

[14] M. G. Haines. A three-dimensional model of wire array instability, ablation, and jetting. *IEEE Transactions on Plasma Science*, 30:588–592, 4 2002. 12, 94

[15] D. B. Sinars, T. A. Shelkovenko, S. A. Pikuz, J. B. Greenly, and D. A. Hammer. Exploding aluminum wire expansion rate with 1–4.5 ka per wire. *Physics of Plasmas*, 7:1555–1563, 5 2000. 12, 95

[16] S. A. Pikuz, T. A. Shelkovenko, D. B. Sinars, J. B. Greenly, Y. S. Dimant, and D. A. Hammer. Multiphase foamlike structure of exploding wire cores. *Physical Review Letters*, 83:4313, 11 1999. 12, 91, 95

[17] J. Ruiz-Camacho, F. N. Beg, A. E. Dangor, M. G. Haines, E. L. Clark, and I. Ross. Z-pinch discharges in aluminum and tungsten wires. *Physics of Plasmas*, 6:2579–2587, 6 1999. 12, 91, 96, 97

[18] J. P. Chittenden, S. V. Lebedev, J. Ruiz-Camacho, F. N. Beg, S. N. Bland, C. A. Jennings, A. R. Bell, M. G. Haines, S. A. Pikuz, T. A. Shelkovenko, and D. A. Hammer. Plasma formation in metallic wire z-pinches. *Physical Review E*, 61:4370, 4 2000. 12, 13, 96, 97, 98, 100

[19] G. W. Collins, M. P. Valdivia, S. B. Hansen, F. Conti, L. C. Carlson, D. A. Hammer, A. Elshafiey, J. Narkis, and F. N. Beg. Direct comparison of wire, foil, and hybrid x-pinches on a 200 ka, 150 ns current driver. *Journal of Applied Physics*, 129, 2 2021. 14, 120

[20] T. A. Shelkovenko, S. A. Pikuz, D. B. Sinars, K. M. Chandler, and D. A. Hammer. X pinch plasma development as a function of wire material and current pulse parameters. *IEEE Transactions on Plasma Science*, 30:567–576, 4 2002. 14, 121, 122

[21] T. A. Shelkovenko, D. B. Sinars, S. A. Pikuz, and D. A. Hammer. Radiographic and spectroscopic studies of x-pinch plasma implosion dynamics and x-ray burst emission characteristics. *Physics of Plasmas*, 8:1305–1318, 4 2001. 14, 120, 123, 124

[22] D. B. Sinars, R. D. McBride, S. A. Pikuz, T. A. Shelkovenko, D. F. Wenger, M. E. Cuneo, E. P. Yu, J. P. Chittenden, E. C. Harding, S. B. Hansen, B. P. Peyton, D. J. Ampleford, and C. A. Jennings. Investigation of high-temperature bright plasma x-ray sources produced in 5-ma x-pinch experiments. *Physical Review Letters*, 109:155002, 10 2012. 15, 139, 140, 141, 142, 143, 155

[23] J.D. Hare. *High energy density magnetic reconnection experiments in colliding carbon plasma flows*. PhD thesis, Imperial College London, 2017. 16, 159, 160

[24] Hantao Ji and William Daughton. Phase diagram for magnetic reconnection in heliophysical, astrophysical, and laboratory plasmas. *Physics of Plasmas*, 18:111207, 11 2011. 16, 159, 160

[25] M. R. Martin, C. E. Seyler, and J. B. Greenly. The role of magnetic field in the transition to streaming ablation in wire arrays. *Physics of Plasmas*, 17, 5 2010. 17, 162

[26] A. J. Harvey-Thompson, S. V. Lebedev, S. N. Bland, J. P. Chittenden, G. N. Hall, A. Marocchino, F. Suzuki-Vidal, S. C. Bott, J. B.A. Palmer, and C. Ning. Quantitative analysis of plasma ablation using inverse wire array z pinches. *Physics of Plasmas*, 16:22701, 2 2009. 17, 162, 163

[27] J. D. Hare, J. MacDonald, S. N. Bland, J. Dranczewski, J. W.D. Halliday, S. V. Lebedev, L. G. Suttle, E. R. Tubman, and W. Rozmus. Two-colour interferometry and thomson scattering measurements of a plasma gun. *Plasma Physics and Controlled Fusion*, 61:085012, 6 2019. 17, 163

[28] R. Paul Drake. *Introduction to High-Energy-Density Physics*, pages 1–17. Springer Berlin Heidelberg, Berlin, Heidelberg, 2006. 19

[29] M. E. Savage, K. R. LeChien, M. R. Lopez, B. S. Stoltzfus, W. A. Stygar, D. S. Artery, J. A. Lott, and P. A. Corcoran. Status of the z pulsed power driver. *Digest of Technical Papers-IEEE International Pulsed Power Conference*, pages 983–990, 2011. 20

[30] R. B. Spielman. Pulsed-power innovations for next-generation, high-current drivers. *IEEE Transactions on Plasma Science*, 50:2621–2627, 9 2022. 21

[31] C. Deeney, M. R. Douglas, R. B. Spielman, T. J. Nash, D. L. Peterson, P. L'Eplattenier, G. A. Chandler, J. F. Seamen, and K. W. Struve. Enhancement of x-ray power from a Z pinch using nested-wire arrays. *Phys. Rev. Lett.*, 81:4883–4886, Nov 1998. 21

[32] M. G. Haines, P. D. Lepell, C. A. Coverdale, B. Jones, C. Deeney, and J. P. Apruzese. Ion viscous heating in a magnetohydrodynamically unstable z pinch at over 2×109 kelvin. *Physical Review Letters*, 96:075003, 2 2006. 21

[33] M. E. Cuneo, R. A. Vesey, G. R. Bennett, D. B. Sinars, W. A. Stygar, E. M. Waisman, J. L. Porter, P. K. Rambo, I. C. Smith, S. V. Lebedev, J. P. Chittenden, D. E. Bliss, T. J. Nash, G. A. Chandler, B. B. Afeyan, E. P. Yu, R. B. Campbell, R. G. Adams, D. L. Hanson, T. A. Mehlhorn, and M. K. Matzen. Progress in symmetric icf capsule implosions and wire-array z-pinch source physics for double-pinch-driven hohlraums*. *Plasma Physics and Controlled Fusion*, 48:R1, 1 2006. 21, 22

[34] John Lindl. Development of the indirect-drive approach to inertial confinement fusion and the target physics basis for ignition and gain. *Physics of Plasmas*, 2:3933–4024, 11 1995. 21

[35] M. J. Edwards and et al. Progress towards ignition on the national ignition facility. *Physics of Plasmas*, 20, 7 2013. 21

[36] A. B. Zylstra and et al. Burning plasma achieved in inertial fusion. *Nature 2022 601:7894*, 601:542–548, 1 2022. 21

[37] M. E. Cuneo and et al. Magnetically driven implosions for inertial confinement fusion at sandia national laboratories. *IEEE Transactions on Plasma Science*, 40:3222–3245, 2012. 22

[38] R. A. Vesey, M. C. Herrmann, R. W. Lemke, M. P. Desjarlais, M. E. Cuneo, W. A. Stygar, G. R. Bennett, R. B. Campbell, P. J. Christenson, T. A. Mehlhorn, J. L. Porter, and S. A. Slutz. Target design for high fusion yield with the double z-pinch-driven hohlraum. *Physics of Plasmas*, 14:56302, 5 2007. 22

[39] D. A. Yager-Elorriaga, M. R. Gomez, D. E. Ruiz, S. A. Slutz, A. J. Harvey-Thompson, C. A. Jennings, P. F. Knapp, P. F. Schmit, M. R. Weis, T. J. Awe, G. A. Chandler, M. Mangan, C. E. Myers, J. R. Fein, B. R. Galloway, M. Geissel, M. E. Glinsky, S. B. Hansen, E. C. Harding, D. C. Lamppa, W. E. Lewis, P. K. Rambo, G. K. Robertson, M. E. Savage, G. A. Shipley, I. C. Smith, J. Schwarz, D. J. Ampleford, K. Beckwith, K. J. Peterson, J. L. Porter, G. A. Rochau, and D. B. Sinars. An overview of magneto-inertial fusion on the z machine at sandia national laboratories. *Nuclear Fusion*, 62:042015, 2 2022. 22, 95

[40] D. E. Ruiz, P. F. Schmit, D. A. Yager-Elorriaga, M. R. Gomez, M. R. Weis, C. A. Jennings, A. J. Harvey-Thompson, P. F. Knapp, S. A. Slutz, D. J. Ampleford, K. Beckwith, and M. K. Matzen. Exploring the parameter space of maglif implosions using similarity scaling. ii. current scaling. *Physics of Plasmas*, 30:32708, 3 2023. 23

[41] DW Scudder. Experiments on high density z-pinches formed from solid deuterium fibers. *Bull. Am. Phys. Soc*, 30(9):1408, 1985. 23

[42] J. D. Sethian, A. E. Robson, K. A. Gerber, and A. W. Desilva. Enhanced stability and neutron production in a dense z-pinch plasma formed from a frozen deuterium fiber. *Physical Review Letters*, 59:892, 8 1987. 23

[43] S. V. Lebedev, R. Aliaga-Rossel, J. P. Chittenden, I. H. Mitchell, A. E. Dangor, M. G. Haines, and J. F. Worley. Coronal plasma behavior of the z pinch produced from carbon and cryogenic deuterium fibers. *Physics of Plasmas*, 5:3366–3372, 9 1998. 23

[44] Malcolm G. Haines. Ion beam formation in an m = 0 unstable z pinch. *Nuclear Instruments and Methods in Physics Research*, 207:179–185, 3 1983. 23

[45] R. S. Pease. Equilibrium characteristics of a pinched gas discharge cooled by bremsstrahlung radiation. *Proceedings of the Physical Society. Section B*, 70:11, 1 1957. 24

[46] L. Bernal and H. Bruzzone. Radiative collapses in z-pinches with axial mass losses. *Plasma Physics and Controlled Fusion*, 44:223, 1 2002. 24

[47] M. G. Haines. The joule heating of a stable pinched plasma. *Proceedings of the Physical Society*, 76:250, 8 1960. 24

[48] J. E. Hammel, D. W. Scudder, and J. S. Shlachter. Recent results on dense z pinches. *Nuclear Instruments and Methods in Physics Research*, 207:161–168, 3 1983. 24

[49] J. W. Shearer. Contraction of z pinches actuated by radiation losses. *The Physics of Fluids*, 19:1426–1428, 9 1976. 24

[50] I.H. Mitchell, P. Choi, J.P. Chittenden, J.F. Worley, J.M. Bayley, A.E. Dangor, and M.G. Haines. The magpie project : Current status. pages 474–479, 1992. 25

[51] J. P. Chittenden and M. G. Haines. Processes terminating radiative collapse in a hydrogen fiber z pinch. *Physics of Fluids B: Plasma Physics*, 2:1889–1897, 8 1990. 25

[52] G.B. Rybicki and A.P. Lightman. *Radiative Processes in Astrophysics*. A Wiley-Interscience publication. Wiley, 1991. 25

[53] D. Griffiths and D. Schroeter. *Introduction to Quantum Mechanics (3rd ed.)*. Cambridge University Press, 2018. 26

[54] A. E. Robson. Lower-hybrid-drift instability and radiative collapse of a dense z pinch. *Physics of Fluids B: Plasma Physics*, 3:1461–1466, 6 1991. 26

[55] J. P. Chittenden. The effect of lower hybrid instabilities on plasma confinement in fiber z pinches. *Physics of Plasmas*, 2:1242–1249, 4 1995. 26

[56] M G Kuzman, P D Agüero, J C Tulli, al, Federico Agustín Membiela, Mauricio Bellini, KJ Montes, C Rea, Ra Tinguely, L Bernal, and H Bruzzone. Radiative collapses in z-pinches with axial mass losses. *Plasma Physics and Controlled Fusion*, 44:223, 1 2002. 27

[57] R. Kingham. Plasma physics postgraduate lecture notes, 2014. 27

[58] M. J. Berger and P. Colella. Local adaptive mesh refinement for shock hydrodynamics. *Journal of Computational Physics*, 82:64–84, 5 1989. 28, 56, 62, 76

[59] E.F. Toro. *Riemann Solvers and Numerical Methods for Fluid Dynamics: A Practical Introduction*. Springer Berlin Heidelberg, 2009. 34

[60] T. D. Arber, A. W. Longbottom, C. L. Gerrard, and A. M. Milne. A staggered grid, lagrangian–eulerian remap code for 3-d mhd simulations. *Journal of Computational Physics*, 171:151–181, 7 2001. 35

[61] Mark L. Wilkins. Use of artificial viscosity in multidimensional fluid dynamic calculations. *Journal of Computational Physics*, 36:281–303, 7 1980. 35

[62] D.L. Youngs. Time-dependent multi-material flow with large fluid distortion. In K. W. Morton and M. J. Baines, editors, *Numerical methods for fluid dynamics*, The Institute of Mathematics and its Applications conference series, pages 273–285. The Institute of Mathematics and its Applications, Academic Press, 1982. 35

[63] Takahiro Miyoshi and Kanya Kusano. A comparative study of divergence-cleaning techniques for multi-dimensional mhd schemes. *Plasma and Fusion Research*, 6:2401124–2401124, 8 2011. 36

[64] Kane S. Yee. Numerical solution of initial boundary value problems involving maxwell's equations in isotropic media. *IEEE Transactions on Antennas and Propagation*, 14:302–307, 1966. 36

[65] C.A. Jennings. *Radiation Transport Effects in Wire Array Z-Pinches and Magneto-Hydrodynamic Modelling Techniques*. PhD thesis, Imperial College London, 2005. 37, 54

[66] Thomas G. White, Jiayu Dai, and David Riley. Dynamic and transient processes in warm dense matter. *Philosophical Transactions of the Royal Society A*, 381, 8 2023. 40

[67] S. Atzeni and J. Meyer-ter-Vehn. *The Physics of Inertial Fusion*. Oxford University Press, 2004. 41

[68] R. M. More, K. H. Warren, D. A. Young, and G. B. Zimmerman. A new quotidian equation of state (QEOS) for hot dense matter. *Phys. Fluids*, 31:3059, 1988. 41, 43, 45

[69] L. H. Thomas. The calculation of atomic fields. *Mathematical Proceedings of the Cambridge Philosophical Society*, 23:542–548, 1927. 41

[70] E. Fermi. Eine statistische methode zur bestimmung einiger eigenschaften des atoms und ihre anwendung auf die theorie des periodischen systems der elemente. *Zeitschrift für Physik*, 48:73–79, 1 1928. 41

[71] E. M. Epperlein and M. G. Haines. Plasma transport coefficients in a magnetic field by direct numerical solution of thr Fokker-Planck equation. *Phys. Fluids*, 29:1029, 1986. 45, 46

[72] L. Spitzer and R. Härm. Transport phenomena in a completely ionized gas. *Phys. Rev.*, 89:977, 1953. 46

[73] John I. Castor. *Radiation Hydrodynamics*. Cambridge University Press, 2004. 49

[74] N.P.L. Niasse. *Development of a Pseudo Non-LTE model for Z-pinch simulations*. PhD thesis, Imperial College London, 2012. 52, 152

[75] A. Crilly, N. Niasse, A. Fraser, D. Chapman, K. Mclean, S. Rose, and J. Chittenden. Spk: A fast atomic and microphysics code for the high-energy-density regime. 2022. 52, 152

[76] J. Huba. *NRL Plasma Formulary*. Naval Research Laboratory, 2016. 52

[77] J.P. Chittenden S. Chowdhry K. Chandler N. Chaturvedi C.E. Myers W.R. Fox S.B. Hansen C.A. Jennings H. Ji C.C. Kuranz S.V. Lebedev D.A. Uzdensky R. Datta, A. Crilly and J.D. Hare. Simulations of radiatively-cooled magnetic reconnection driven by pulsed power. unpublished, 2024. 53, 165, 168, 174, 176

[78] Hesham El-Rewini and Mostafa Abd-El-Barr. Advanced computer architecture and parallel processing. *Advanced Computer Architecture and Parallel Processing*, pages 1–272, 1 2005. 58, 59

[79] G. Hager and G. Wellein. *Introduction to High Performance Computing for Scientists and Engineers*. Chapman & Hall/CRC Computational Science. CRC Press, 2010. 58

[80] Michael Griebel and Gerhard Zumbusch. Parallel multigrid in an adaptive pde solver based on hashing and space-filling curves. *Parallel Computing*, 25(7):827 – 843, 1999. 61

[81] A. M. Khokhlov. Fully threaded tree algorithms for adaptive refinement fluid dynamics simulations. *Journal of Computational Physics*, 143:519–543, 7 1998. 61

[82] Tomoya Ogawa, Takuma Ohta, Ryoji Matsumoto, Kazuyuki Yamashita, and Mitsue Den. Hydrodynamical simulations using a fully threaded tree. *Progress of Theoretical Physics Supplement*, 138:654–655, 4 2000. 61

[83] Teyssier, R. Cosmological hydrodynamics with adaptive mesh refinement - a new high resolution code called ramses. *A&A*, 385(1):337–364, 2002. 61

[84] Weiqun Zhang, Ann Almgren, Vince Beckner, John Bell, Johannes Blaschke, Cy Chan, Marcus Day, Brian Friesen, Kevin Gott, Daniel Graves, Max P. Katz, Andrew Myers, Tan Nguyen, Andrew Nonaka, Michele Rosso, Samuel Williams, and Michael Zingale. Amrex: a framework for block-structured adaptive mesh refinement. *Journal of Open Source Software*, 4(37):1370, 2019. 62

[85] W. Zhang, A. Myers, K. Gott, A. Almgren, and J. Bell. Amrex: Block-structured adaptivemesh refinement for multiphysics applications. *ArXiv*, abs/2009.12009, September 2020. 62

[86] B. Fryxell, K. Olson, P. Ricker, F. X. Timmes, M. Zingale, D. Q. Lamb, P. MacNeice, R. Rosner, J. W. Truran, and H. Tufo. Flash: An adaptive mesh hydrodynamics code for modeling astrophysical thermonuclear flashes. *The Astrophysical Journal Supplement Series*, 131(1):273, nov 2000. 63, 81

[87] Greg L. Bryan, Michael L. Norman, Brian W. O'Shea, Tom Abel, John H. Wise, Matthew J. Turk, Daniel R. Reynolds, David C. Collins, Peng Wang, Samuel W. Skillman, Britton Smith, Robert P. Harkness, James Bordner, Ji hoon Kim, Michael Kuhlen, Hao Xu, Nathan Goldbaum, Cameron Hummels, Alexei G. Kritsuk, Elizabeth Tasker, Stephen Skory, Christine M. Simpson, Oliver Hahn, Jeffrey S. Oishi, Geoffrey C. So, Fen Zhao, Renyue Cen, Yuan Li, and (The Enzo Collaboration). Enzo: An adaptive mesh refinement code for astrophysics. *The Astrophysical Journal Supplement Series*, 211(2):19, mar 2014. 63, 76, 81

[88] James M. Stone, Kengo Tomida, Christopher J. White, and Kyle G. Felker. The athena++ adaptive mesh refinement framework: Design and magnetohydrodynamic solvers. *The Astrophysical Journal Supplement Series*, 249(1):4, jun 2020. 63

[89] G.M. Morton. *A Computer Oriented Geodetic Data Base and a New Technique in File Sequencing*. International Business Machines Company, 1966. 63

[90] Will Schroeder, Ken Martin, and Bill Lorensen. *The Visualization Toolkit (4th ed.)*. Kitware, 2006. 78

[91] James Ahrens, Berk Geveci, and Charles Law. *Visualization Handbook*, chapter ParaView: An End-User Tool for Large Data Visualization, pages 717–731. Elsevier Inc., Burlington, MA, USA, 2005. 78

[92] Charles R. Evans and John F. Hawley. Simulation of Magnetohydrodynamic Flows: A Constrained Transport Model. , 332:659, Sep 1988. 80

[93] Dinshaw S. Balsara, Roger Käppeli, Walter Boscheri, and Michael Dumbser. Curl constraint-preserving reconstruction and the guidance it gives for mimetic scheme design. *Communications on Applied Mathematics and Computation*, 5:235–294, 3 2023. 81

[94] D. H. Kalantar and D. A. Hammer. Observation of a stable dense core within an unstable coronal plasma in wire-initiated dense z-pinch experiments. *Physical Review Letters*, 71:3806, 12 1993. 91

[95] F. N. Beg, A. E. Dangor, P. Lee, M. Tatarakis, S. L. Niffikeer, and M. G. Haines. Optical and x-ray observations of carbon and aluminium fibre z-pinch plasmas. *Plasma Physics and Controlled Fusion*, 39:1, 1 1997. 91

[96] G. V. Ivanenkov, A. R. Mingaleev, S. A. Pikuz, V. M. Romanova, T. A. Shelkovenko, W. Stepniewski, and D. A. Hammer. Dynamics of thin exploded-wire plasma with a cold dense core. *Journal of Experimental and Theoretical Physics*, 87:663–671, 1998. 91

[97] G. S. Sarkisov, P. V. Sasorov, K. W. Struve, D. H. McDaniel, A. N. Gribov, and G. M. Oleinik. Polarity effect for exploding wires in a vacuum. *Physical Review E*, 66:046413, 10 2002. 92

[98] S N Bland, S V Lebedev, J P Chittenden, D J Ampleford, S C Bott, J A Gómez, M G Haines, G N Hall, D A Hammer, I H Mitchell, and J B A Palmer. Effect of radial-electric-field polarity on wire-array z-pinch dynamics. 2005. 92

[99] S. V. Lebedev, F. N. Beg, S. N. Bland, J. P. Chittenden, A. E. Dangor, M. G. Haines, S. A. Pikuz, and T. A. Shelkovenko. Effect of core-corona plasma structure on seeding of instabilities in wire array z-pinches. *Phys. Rev. Lett.*, 85:98–101, Jul 2000. 93

[100] G. N. Hall, J. P. Chittenden, S. N. Bland, S. V. Lebedev, S. C. Bott, C. Jennings, J. B.A. Palmer, and F. Suzuki-Vidal. Modifying wire-array z-pinch ablation structure using coiled arrays. *Physical Review Letters*, 100:065003, 2 2008. 93

[101] J. P. Chittenden and C. A. Jennings. Development of instabilities in wire-array z pinches. *Physical Review Letters*, 101:055005, 8 2008. 94, 112

[102] E. P. Yu, T. J. Awe, K. R. Cochrane, K. J. Peterson, K. C. Yates, T. M. Hutchinson, M. W. Hatch, B. S. Bauer, K. Tomlinson, and D. B. Sinars. Seeding the electrothermal instability through a three-dimensional, nonlinear perturbation. *Physical Review Letters*, 130:255101, 6 2023. 95

[103] I. H. Mitchell, J. M. Bayley, J. P. Chittenden, J. F. Worley, A. E. Dangor, M. G. Haines, and P. Choi. A high impedance mega-ampere generator for fiber z-pinch experiments. *Review of Scientific Instruments*, 67:1533–1541, 4 1996. 96

[104] SM Zakharov, GV Ivanenkov, AA Kolomenskij, SA Pikuz, AI Samokhin, and I Ulshmid. Wire x-pinch in a high-current diode. *Pis' ma v Zhurnal Tekhnicheskoj Fiziki*, 8(17):1060–1063, 1982. 119

[105] S. A. Pikuz, T. A. Shelkovenko, and D. A. Hammer. X-pinch. part ii. *Plasma Physics Reports*, 41:445–491, 6 2015. 122

[106] Philip A Hughes. *Beams and jets in astrophysics*. Number 19. Cambridge University Press, 1991. 122

[107] S. V. Lebedev, J. P. Chittenden, F. N. Beg, S. N. Bland, A. Ciardi, D. Ampleford, S. Hughes, M. G. Haines, A. Frank, E. G. Blackman, and T. Gardiner. Laboratory astrophysics and collimated stellar outflows: The production of radiatively cooled hypersonic plasma jets. *The Astrophysical Journal*, 564:113–119, 1 2002. 122

[108] Andrea L. Kritcher, Damian C. Swift, Tilo Döppner, Benjamin Bachmann, Lorin X. Benedict, Gilbert W. Collins, Jonathan L. DuBois, Fred Elsner, Gilles Fontaine, Jim A. Gaffney, Sebastien Hamel, Amy Lazicki, Walter R. Johnson, Natalie Kostinski, Dominik Kraus, Michael J. MacDonald, Brian Maddox, Madison E. Martin, Paul Neumayer, Abbas Nikroo, Joseph Nilsen, Bruce A. Remington, Didier Saumon, Phillip A. Sterne, Wendi Sweet, Alfredo A. Correa, Heather D. Whitley, Roger W. Falcone, and Siegfried H. Glenzer. A measurement of the equation of state of carbon envelopes of white dwarfs. *Nature 2020 584:7819*, 584:51–54, 8 2020. 122

[109] K. C. Mittal, D. H. Kalantar, N. Qi, D. A. Hammer, K. A. Gerber, and J. D. Sethian. Neutron production in dense x-pinch plasmas produced from deuterated polyethylene fibers. *Journal of Applied Physics*, 70:6712–6717, 12 1991. 122

[110] S. S. Anan'ev, Yu. L. Bakshaev, P. I. Blinov, V. A. Bryzgunov, V. V. Vikhrev, S. A. Dan'ko, A. A. Zelenin, E. D. Kazakov, Yu. G. Kalinin, A. S. Kingsep, V. D. Korolev, E. A. Smirnova, G. I. Ustroev, A. S. Chernenko, and V. A. Shchagin. X-pinch-based neutron source. *Plasma Physics Reports*, 36(7):601–608, July 2010. 122

[111] J. D. Douglass, J. B. Greenly, D. A. Hammer, R. D. McBride, S. A. Pikuz, and T. A. Shelkovenko. The imaging of z-pinches using x-pinch backlighting. *AIP Conference Proceedings*, 808:129–132, 1 2006. 122

[112] Clayton Edward Myers, Matthew R Gomez, Derek C. Lamppa, Timothy Jay Webb, David Alexander Yager-Elorriaga, Brian Thomas Hutsel, Christopher Ashley Jennings, Patrick Knapp, Michael Richard Kossow, Larry Martin Lucero,

Robert J. Obregon, and Adam Michael Steiner. First demonstration of an inductively driven x-pinch for diagnosing high energy density experiments on the z pulsed power facility. 11 2020. 122

[113] J. P. Chittenden, A. Ciardi, C. A. Jennings, S. V. Lebedev, D. A. Hammer, S. A. Pikuz, and T. A. Shelkovenko. Structural evolution and formation of high-pressure plasmas in x pinches. *Physical Review Letters*, 98:025003, 1 2007. 131

[114] A. Skoulakis, G. Koundourakis, A. Ciardi, E. Kaselouris, I. Fitilis, J. Chatzakis, M. Bakarezos, N. Vlahakis, N. A. Papadogiannis, M. Tatarakis, and V. Dimitriou. High performance simulations of a single x-pinch. *Plasma Physics and Controlled Fusion*, 64:025003, 12 2021. 137

[115] D. D. Ryutov, M. S. Derzon, and M. K. Matzen. The physics of fast z-pinches. *Reviews of Modern Physics*, 72:167, 1 2000. 139

[116] D. B. Sinars, S. A. Pikuz, J. D. Douglass, R. D. McBride, D. J. Ampleford, P. Knapp, K. Bell, D. Chalenksi, M. E. Cuneo, J. B. Greenly, D. A. Hammer, B. R. Kusse, A. Mingaleev, T. A. Shelkovenko, and D. F. Wenger. Bright spots in 1 ma x pinches as a function of wire number and material. *Physics of Plasmas*, 15:92703, 9 2008. 139

[117] D. Wenger D. Sinars J. Chittenden S. Pikuz E. Harding C.A. Jennings D. Ampleford E. Yu M. Cuneo T Shelkovenko S. Hansen S. Bland, R. McBride. Scaling of x pinches from 1 ma to 6 ma. 9 2010. 140, 141

[118] Byung Moo Song, Sergei A. Pikuz, Tatiania A. Shelkovenko, and David A. Hammer. Determination of the size and structure of an x-pinch x-ray source from the diffraction pattern produced by microfabricated slits. *Applied Optics, Vol. 44, Issue 12, pp. 2349-2358*, 44:2349–2358, 4 2005. 140

[119] S. P. Regan and et al. Demonstration of fuel hot-spot pressure in excess of 50 gbar for direct-drive, layered deuterium-tritium implosions on omega. *Phys. Rev. Lett.*, 117:025001, Jul 2016. 147

[120] A. Kramida, Yu. Ralchenko, J. Reader, and and NIST ASD Team. NIST Atomic Spectra Database (ver. 5.10), [Online]. Available: `https://physics.nist.gov/asd` [2023, November 5]. National Institute of Standards and Technology, Gaithersburg, MD., 2022. 155

[121] D. A. Uzdensky. *Radiative Magnetic Reconnection in Astrophysics*, pages 473–519. Springer International Publishing, Cham, 2016. 158

[122] M. Hesse and P. A. Cassak. Magnetic reconnection in the space sciences: Past, present, and future. *Journal of Geophysical Research: Space Physics*, 125:e2018JA025935, 2 2020. 159

[123] P. A. Sweet. The neutral point theory of solar flares. *Symposium - International Astronomical Union*, 6:123–134, 1958. 159

[124] Eugene N Parker. Sweet's mechanism for merging magnetic fields in conducting fluids. *Journal of Geophysical Research*, 62(4):509–520, 1957. 159

[125] N. Loureiro and D Uzdensky. Magnetic reconnection: From the sweet-parker model to stochastic plasmoid chains. *Plasma Physics and Controlled Fusion*, 58, 08 2015. 159

[126] Dmitri A. Uzdensky and Jonathan C. McKinney. Magnetic reconnection with radiative cooling. i. optically thin regime. *Physics of Plasmas*, 18:42105, 4 2011. 160

[127] Masaaki Yamada, Hantao ji, Scott Hsu, Troy Carter, Russell Kulsrud, Norton Bretz, Forrest Jobes, Yasushi Ono, and Francis Perkins. Study of driven magnetic reconnection in a laboratory plasma. *Physics of Plasmas*, 4:1936–1944, 5 1997. 161

[128] P. M. Nilson, L. Willingale, M. C. Kaluza, C. Kamperidis, S. Minardi, M. S. Wei, P. Fernandes, M. Notley, S. Bandyopadhyay, M. Sherlock, R. J. Kingham, M. Tatarakis, Z. Najmudin, W. Rozmus, R. G. Evans, M. G. Haines, A. E. Dangor, and K. Krushelnick. Magnetic reconnection and plasma dynamics in two-beam laser-solid interactions. *Physical Review Letters*, 97:255001, 12 2006. 161

[129] J. D. Hare, L. G. Suttle, S. V. Lebedev, N. F. Loureiro, A. Ciardi, J. P. Chittenden, T. Clayson, S. J. Eardley, C. Garcia, J. W.D. Halliday, T. Robinson, R. A. Smith, N. Stuart, F. Suzuki-Vidal, and E. R. Tubman. An experimental platform for pulsed-power driven magnetic reconnection. *Physics of Plasmas*, 25:55703, 5 2018. 161, 163, 164

[130] D.R. Russell. *Bow shock interaction experiments in a magnetised collisional plasma.* PhD thesis, Imperial College London, 2021. 163

[131] G. C. Burdiak, S. V. Lebedev, A. J. Harvey-Thompson, G. N. Hall, G. F. Swadling, F. Suzuki-Vidal, E. Khoory, S. N. Bland, L. Pickworth, P. De Grouchy, J. Skidmore, L. Suttle, and E. M. Waisman. Characterisation of the current switch mechanism in two-stage wire array z-pinches. *Physics of Plasmas*, 22, 11 2015. 163

[132] L. G. Suttle, J. D. Hare, J. W.D. Halliday, S. Merlini, D. R. Russell, E. R. Tubman, V. Valenzuela-Villaseca, W. Rozmus, C. Bruulsema, and S. V. Lebedev. Collective optical thomson scattering in pulsed-power driven high energy density physics experiments (invited). *Review of Scientific Instruments*, 92:33542, 3 2021. 163

[133] D. B. Sinars and et al. Review of pulsed power-driven high energy density physics research on Z at Sandia. *Physics of Plasmas*, 27(7):070501, 07 2020. 167

Figure permissions

AIP PUBLISHING LICENSE
TERMS AND CONDITIONS

Dec 30, 2023

This Agreement between Ms. Nikita Chaturvedi ("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	5698900118260
License date	Dec 30, 2023
Licensed Content Publisher	AIP Publishing
Licensed Content Publication	Physics of Plasmas
Licensed Content Title	Steady-state radiation ablation in the wire-array Z pinch
Licensed Content Author	Yu, Edmund P.; Oliver, B. V.
Licensed Content Date	Feb 20, 2007
Licensed Content Volume	14
Licensed Content Issue	2
Type of Use	Thesis/Dissertation
Requestor type	Student
Format	Print
Portion	Figure/Table

Number of figures/tables	1
Will you be translating?	No
Title of new work	Simulating radiatively cooled pulsed-power driven systems using static mesh refinement
Institution name	Imperial College London
Expected presentation date	Dec 2023
Portions	Fig 5
Requestor Location	Ms. Nikita Chaturvedi 28 Westminster Palace Gardens London, SW1P 1RR United Kingdom Attn: Ms. Nikita Chaturvedi
Total	0.00 USD

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.

American Physical Society
Reuse and Permissions License

15-Feb-2024

This license agreement between the American Physical Society ("APS") and Nikita Chaturvedi ("You") consists of your license details and the terms and conditions provided by the American Physical Society and SciPris.

Licensed Content Information

License Number:	**RNP/24/FEB/075137**
License date:	15-Feb-2024
DOI:	10.1103/PhysRevLett.83.4313
Title:	Multiphase Foamlike Structure of Exploding Wire Cores
Author:	S. A. Pikuz et al.
Publication:	Physical Review Letters
Publisher:	American Physical Society
Cost:	USD $ 0.00

Request Details

Does your reuse require significant modifications: No

Specify intended distribution locations:	Worldwide
Reuse Category:	Reuse in a thesis/dissertation
Requestor Type:	Student
Items for Reuse:	Figures/Tables
Number of Figure/Tables:	1
Figure/Tables Details:	Fig 1.3
Format for Reuse:	Electronic

Information about New Publication:

University/Publisher:	Imperial College London
Title of dissertation/thesis:	Simulating radiatively cooled pulsed-power driven systems using static mesh refinement
Author(s):	Nikita Chaturvedi
Expected completion date:	Dec. 2023

License Requestor Information

Name:	Nikita Chaturvedi
Affiliation:	Individual
Email Id:	nc1613@ic.ac.uk
Country:	United Kingdom

American Physical Society
Reuse and Permissions License

TERMS AND CONDITIONS

The American Physical Society (APS) is pleased to grant the Requestor of this license a non-exclusive, non-transferable permission, limited to Electronic format, provided all criteria outlined below are followed.

1. You must also obtain permission from at least one of the lead authors for each separate work, if you haven't done so already. The author's name and affiliation can be found on the first page of the published Article.
2. For electronic format permissions, Requestor agrees to provide a hyperlink from the reprinted APS material using the source material's DOI on the web page where the work appears. The hyperlink should use the standard DOI resolution URL, http://dx.doi.org/{DOI}. The hyperlink may be embedded in the copyright credit line.
3. For print format permissions, Requestor agrees to print the required copyright credit line on the first page where the material appears: "Reprinted (abstract/excerpt/figure) with permission from [(FULL REFERENCE CITATION) as follows: Author's Names, APS Journal Title, Volume Number, Page Number and Year of Publication.] Copyright (YEAR) by the American Physical Society."
4. Permission granted in this license is for a one-time use and does not include permission for any future editions, updates, databases, formats or other matters. Permission must be sought for any additional use.
5. Use of the material does not and must not imply any endorsement by APS.
6. APS does not imply, purport or intend to grant permission to reuse materials to which it does not hold copyright. It is the requestor's sole responsibility to ensure the licensed material is original to APS and does not contain the copyright of another entity, and that the copyright notice of the figure, photograph, cover or table does not indicate it was reprinted by APS with permission from another source.
7. The permission granted herein is personal to the Requestor for the use specified and is not transferable or assignable without express written permission of APS. This license may not be amended except in writing by APS.
8. You may not alter, edit or modify the material in any manner.
9. You may translate the materials only when translation rights have been granted.
10. APS is not responsible for any errors or omissions due to translation.
11. You may not use the material for promotional, sales, advertising or marketing purposes.
12. The foregoing license shall not take effect unless and until APS or its agent, Aptara, receives payment in full in accordance with Aptara Billing and Payment Terms and Conditions, which are incorporated herein by reference.
13. Should the terms of this license be violated at any time, APS or Aptara may revoke the license with no refund to you and seek relief to the fullest extent of the laws of the USA. Official written notice will be made using the contact information provided with the permission request. Failure to receive such notice will not nullify revocation of the permission.
14. APS reserves all rights not specifically granted herein.
15. This document, including the Aptara Billing and Payment Terms and Conditions, shall be the entire agreement between the parties relating to the subject matter hereof.

AIP PUBLISHING LICENSE
TERMS AND CONDITIONS

Feb 15, 2024

This Agreement between Ms. Nikita Chaturvedi ("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	5730391036527
License date	Feb 15, 2024
Licensed Content Publisher	AIP Publishing
Licensed Content Publication	Physics of Plasmas
Licensed Content Title	Z-pinch discharges in aluminum and tungsten wires
Licensed Content Author	Ruiz-Camacho, J.; Beg, F. N.
Licensed Content Date	Jun 1, 1999
Licensed Content Volume	6
Licensed Content Issue	6
Type of Use	Thesis/Dissertation
Requestor type	Student
Format	Print and electronic
Portion	Figure/Table

Number of figures/tables	2
Will you be translating?	No
Title of new work	Simulating radiatively cooled pulsed-power driven systems using static mesh refinement
Institution name	Imperial College London
Expected presentation date	Feb 2024
Portions	6,7
Requestor Location	Ms. Nikita Chaturvedi 28 Westminster Palace Gardens London, SW1P 1RR United Kingdom Attn: Ms. Nikita Chaturvedi
Total	0.00 USD

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.

AIP PUBLISHING LICENSE
TERMS AND CONDITIONS

Feb 15, 2024

This Agreement between Ms. Nikita Chaturvedi ("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	5730390931384
License date	Feb 15, 2024
Licensed Content Publisher	AIP Publishing
Licensed Content Publication	Journal of Applied Physics
Licensed Content Title	Direct comparison of wire, foil, and hybrid X-pinches on a 200 kA, 150 ns current driver
Licensed Content Author	Collins, G. W.; Valdivia, M. P.
Licensed Content Date	Feb 18, 2021
Licensed Content Volume	129
Licensed Content Issue	7
Type of Use	Thesis/Dissertation
Requestor type	Student
Format	Print and electronic
Portion	Figure/Table

Number of figures/tables	1
Will you be translating?	No
Title of new work	Simulating radiatively cooled pulsed-power driven systems using static mesh refinement
Institution name	Imperial College London
Expected presentation date	Feb 2024
Portions	Fig 1
Requestor Location	Ms. Nikita Chaturvedi 28 Westminster Palace Gardens London, SW1P 1RR United Kingdom Attn: Ms. Nikita Chaturvedi
Total	0.00 USD

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.

AIP PUBLISHING LICENSE
TERMS AND CONDITIONS

Feb 15, 2024

This Agreement between Ms. Nikita Chaturvedi ("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	5730390854282
License date	Feb 15, 2024
Licensed Content Publisher	AIP Publishing
Licensed Content Publication	Physics of Plasmas
Licensed Content Title	Radiographic and spectroscopic studies of X-pinch plasma implosion dynamics and x-ray burst emission characteristics
Licensed Content Author	Shelkovenko, T. A.; Sinars, D. B.
Licensed Content Date	Apr 1, 2001
Licensed Content Volume	8
Licensed Content Issue	4
Type of Use	Thesis/Dissertation
Requestor type	Student
Format	Print and electronic

Portion	Figure/Table
Number of figures/tables	1
Will you be translating?	No
Title of new work	Simulating radiatively cooled pulsed-power driven systems using static mesh refinement
Institution name	Imperial College London
Expected presentation date	Feb 2024
Portions	7
Requestor Location	Ms. Nikita Chaturvedi 28 Westminster Palace Gardens London, SW1P 1RR United Kingdom Attn: Ms. Nikita Chaturvedi
Total	0.00 USD

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.

American Physical Society
Reuse and Permissions License

29-Oct-2023

This license agreement between the American Physical Society ("APS") and Nikita Chaturvedi ("You") consists of your license details and the terms and conditions provided by the American Physical Society and SciPris.

Licensed Content Information

License Number:	**RNP/23/OCT/071481**
License date:	29-Oct-2023
DOI:	10.1103/PhysRevLett.109.155002
Title:	Investigation of High-Temperature Bright Plasma X-ray Sources Produced in 5-MA X-Pinch Experiments
Author:	D. B. Sinars et al.
Publication:	Physical Review Letters
Publisher:	American Physical Society
Cost:	USD $ 0.00

Request Details

Does your reuse require significant modifications: No

Specify intended distribution locations:	Worldwide
Reuse Category:	Reuse in a journal/magazine
Requestor Type:	Academic Institution
Items for Reuse:	Figures/Tables
Number of Figure/Tables:	3
Figure/Tables Details:	Figs 1,2, and 4 for comparison against simulation results
Format for Reuse:	Electronic

Information about New Publication:

Publisher:	IEEE
Publication:	IEEE Transactions on Plasma Science
Publication Date:	May. 2024
Article Title:	Rad-MHD simulations of multi-MA X-pinches
Author(s):	N. Chaturvedi, J.P. Chittenden

License Requestor Information

Name:	Nikita Chaturvedi
Affiliation:	Individual
Email Id:	nc1613@ic.ac.uk
Country:	United Kingdom

American Physical Society
Reuse and Permissions License

TERMS AND CONDITIONS

The American Physical Society (APS) is pleased to grant the Requestor of this license a non-exclusive, non-transferable permission, limited to Electronic format, provided all criteria outlined below are followed.

1. You must also obtain permission from at least one of the lead authors for each separate work, if you haven't done so already. The author's name and affiliation can be found on the first page of the published Article.
2. For electronic format permissions, Requestor agrees to provide a hyperlink from the reprinted APS material using the source material's DOI on the web page where the work appears. The hyperlink should use the standard DOI resolution URL, http://dx.doi.org/{DOI}. The hyperlink may be embedded in the copyright credit line.
3. For print format permissions, Requestor agrees to print the required copyright credit line on the first page where the material appears: "Reprinted (abstract/excerpt/figure) with permission from [(FULL REFERENCE CITATION) as follows: Author's Names, APS Journal Title, Volume Number, Page Number and Year of Publication.] Copyright (YEAR) by the American Physical Society."
4. Permission granted in this license is for a one-time use and does not include permission for any future editions, updates, databases, formats or other matters. Permission must be sought for any additional use.
5. Use of the material does not and must not imply any endorsement by APS.
6. APS does not imply, purport or intend to grant permission to reuse materials to which it does not hold copyright. It is the requestor's sole responsibility to ensure the licensed material is original to APS and does not contain the copyright of another entity, and that the copyright notice of the figure, photograph, cover or table does not indicate it was reprinted by APS with permission from another source.
7. The permission granted herein is personal to the Requestor for the use specified and is not transferable or assignable without express written permission of APS. This license may not be amended except in writing by APS.
8. You may not alter, edit or modify the material in any manner.
9. You may translate the materials only when translation rights have been granted.
10. APS is not responsible for any errors or omissions due to translation.
11. You may not use the material for promotional, sales, advertising or marketing purposes.
12. The foregoing license shall not take effect unless and until APS or its agent, Aptara, receives payment in full in accordance with Aptara Billing and Payment Terms and Conditions, which are incorporated herein by reference.
13. Should the terms of this license be violated at any time, APS or Aptara may revoke the license with no refund to you and seek relief to the fullest extent of the laws of the USA. Official written notice will be made using the contact information provided with the permission request. Failure to receive such notice will not nullify revocation of the permission.
14. APS reserves all rights not specifically granted herein.
15. This document, including the Aptara Billing and Payment Terms and Conditions, shall be the entire agreement between the parties relating to the subject matter hereof.

AIP PUBLISHING LICENSE TERMS AND CONDITIONS

Feb 15, 2024

This Agreement between Ms. Nikita Chaturvedi ("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	5730390537488
License date	Feb 15, 2024
Licensed Content Publisher	AIP Publishing
Licensed Content Publication	Physics of Plasmas
Licensed Content Title	Phase diagram for magnetic reconnection in heliophysical, astrophysical, and laboratory plasmas
Licensed Content Author	Ji, Hantao; Daughton, William
Licensed Content Date	Oct 18, 2011
Licensed Content Volume	18
Licensed Content Issue	11
Type of Use	Thesis/Dissertation
Requestor type	Student
Format	Print and electronic
Portion	Figure/Table
Number of figures/tables	1
Will you be translating?	No
Title of new work	Simulating radiatively cooled pulsed-power driven systems using static mesh refinement
Institution name	Imperial College London
Expected presentation date	Feb 2024
Portions	Fig 1
Requestor Location	Ms. Nikita Chaturvedi 28 Westminster Palace Gardens London, SW1P 1RR United Kingdom Attn: Ms. Nikita Chaturvedi
Total	**0.00 USD**
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? E-mail us at customercare@copyright.com.

AIP PUBLISHING LICENSE
TERMS AND CONDITIONS

Feb 15, 2024

This Agreement between Ms. Nikita Chaturvedi ("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	5730390648867
License date	Feb 15, 2024
Licensed Content Publisher	AIP Publishing
Licensed Content Publication	Physics of Plasmas
Licensed Content Title	The role of magnetic field in the transition to streaming ablation in wire arrays
Licensed Content Author	Martin, M. R.; Seyler, C. E.
Licensed Content Date	May 18, 2010
Licensed Content Volume	17
Licensed Content Issue	5
Type of Use	Thesis/Dissertation
Requestor type	Student
Format	Print and electronic
Portion	Figure/Table

Number of figures/tables	1
Will you be translating?	No
Title of new work	Simulating radiatively cooled pulsed-power driven systems using static mesh refinement
Institution name	Imperial College London
Expected presentation date	Feb 2024
Portions	1
Requestor Location	Ms. Nikita Chaturvedi 28 Westminster Palace Gardens London, SW1P 1RR United Kingdom Attn: Ms. Nikita Chaturvedi
Total	0.00 USD

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.

American Physical Society
Reuse and Permissions License

30-Dec-2023

This license agreement between the American Physical Society ("APS") and Nikita Chaturvedi ("You") consists of your license details and the terms and conditions provided by the American Physical Society and SciPris.

Licensed Content Information

License Number:	**RNP/23/DEC/073492**
License date:	30-Dec-2023
DOI:	10.1103/PhysRevLett.113.155003
Title:	Experimental Demonstration of Fusion-Relevant Conditions in Magnetized Liner Inertial Fusion
Author:	M. R. Gomez et al.
Publication:	Physical Review Letters
Publisher:	American Physical Society
Cost:	USD $ 0.00

Request Details

Does your reuse require significant modifications: No

Specify intended distribution locations:	Worldwide
Reuse Category:	Reuse in a thesis/dissertation
Requestor Type:	Student
Items for Reuse:	Figures/Tables
Number of Figure/Tables:	1
Figure/Tables Details:	Fig 1
Format for Reuse:	Electronic

Information about New Publication:

University/Publisher:	Imperial College London
Title of dissertation/thesis:	Simulating radiatively cooled pulsed-power driven systems using static mesh refinement
Author(s):	Nikita Chaturvedi
Expected completion date:	Dec. 2023

License Requestor Information

Name:	Nikita Chaturvedi
Affiliation:	Individual
Email Id:	nc1613@ic.ac.uk
Country:	United Kingdom

American Physical Society
Reuse and Permissions License

TERMS AND CONDITIONS

The American Physical Society (APS) is pleased to grant the Requestor of this license a non-exclusive, non-transferable permission, limited to Electronic format, provided all criteria outlined below are followed.

1. You must also obtain permission from at least one of the lead authors for each separate work, if you haven't done so already. The author's name and affiliation can be found on the first page of the published Article.
2. For electronic format permissions, Requestor agrees to provide a hyperlink from the reprinted APS material using the source material's DOI on the web page where the work appears. The hyperlink should use the standard DOI resolution URL, http://dx.doi.org/{DOI}. The hyperlink may be embedded in the copyright credit line.
3. For print format permissions, Requestor agrees to print the required copyright credit line on the first page where the material appears: "Reprinted (abstract/excerpt/figure) with permission from [(FULL REFERENCE CITATION) as follows: Author's Names, APS Journal Title, Volume Number, Page Number and Year of Publication.] Copyright (YEAR) by the American Physical Society."
4. Permission granted in this license is for a one-time use and does not include permission for any future editions, updates, databases, formats or other matters. Permission must be sought for any additional use.
5. Use of the material does not and must not imply any endorsement by APS.
6. APS does not imply, purport or intend to grant permission to reuse materials to which it does not hold copyright. It is the requestor's sole responsibility to ensure the licensed material is original to APS and does not contain the copyright of another entity, and that the copyright notice of the figure, photograph, cover or table does not indicate it was reprinted by APS with permission from another source.
7. The permission granted herein is personal to the Requestor for the use specified and is not transferable or assignable without express written permission of APS. This license may not be amended except in writing by APS.
8. You may not alter, edit or modify the material in any manner.
9. You may translate the materials only when translation rights have been granted.
10. APS is not responsible for any errors or omissions due to translation.
11. You may not use the material for promotional, sales, advertising or marketing purposes.
12. The foregoing license shall not take effect unless and until APS or its agent, Aptara, receives payment in full in accordance with Aptara Billing and Payment Terms and Conditions, which are incorporated herein by reference.
13. Should the terms of this license be violated at any time, APS or Aptara may revoke the license with no refund to you and seek relief to the fullest extent of the laws of the USA. Official written notice will be made using the contact information provided with the permission request. Failure to receive such notice will not nullify revocation of the permission.
14. APS reserves all rights not specifically granted herein.
15. This document, including the Aptara Billing and Payment Terms and Conditions, shall be the entire agreement between the parties relating to the subject matter hereof.

ELSEVIER LICENSE
TERMS AND CONDITIONS

Dec 30, 2023

This Agreement between Ms. Nikita Chaturvedi ("You") and Elsevier ("Elsevier") consists of your license details and the terms and conditions provided by Elsevier and Copyright Clearance Center.

License Number	5698890947979
License date	Dec 30, 2023
Licensed Content Publisher	Elsevier
Licensed Content Publication	Computer Physics Communications
Licensed Content Title	The equation of state package FEOS for high energy density matter
Licensed Content Author	Steffen Faik,Anna Tauschwitz,Igor Iosilevskiy
Licensed Content Date	Jun 1, 2018
Licensed Content Volume	227
Licensed Content Issue	n/a
Licensed Content Pages	9
Start Page	117
End Page	125
Type of Use	reuse in a thesis/dissertation
Portion	figures/tables/illustrations

Number of figures/tables/illustrations	1
Format	print
Are you the author of this Elsevier article?	No
Will you be translating?	No
Title of new work	Simulating radiatively cooled pulsed-power driven systems using static mesh refinement
Institution name	Imperial College London
Expected presentation date	Dec 2023
Portions	Fig 3
Requestor Location	Ms. Nikita Chaturvedi 28 Westminster Palace Gardens London, SW1P 1RR United Kingdom Attn: Ms. Nikita Chaturvedi
Publisher Tax ID	GB 494 6272 12
Total	0.00 USD
Terms and Conditions	

INTRODUCTION

1. The publisher for this copyrighted material is Elsevier. By clicking "accept" in connection with completing this licensing transaction, you agree that the following terms and conditions apply to this transaction (along with the Billing and Payment terms and conditions established by Copyright Clearance Center, Inc. ("CCC"), at the time that you opened your RightsLink account and that are available at any time at https://myaccount.copyright.com).

GENERAL TERMS

2. Elsevier hereby grants you permission to reproduce the aforementioned material subject to the terms and conditions indicated.

3. Acknowledgement: If any part of the material to be used (for example, figures) has appeared in our publication with credit or acknowledgement to another source, permission must also be sought from that source. If such permission is not obtained then that material may not be included in your publication/copies. Suitable acknowledgement to the source must be made, either as a footnote or in a reference list at the end of your publication, as follows:

"Reprinted from Publication title, Vol /edition number, Author(s), Title of article / title of chapter, Pages No., Copyright (Year), with permission from Elsevier [OR APPLICABLE SOCIETY COPYRIGHT OWNER]." Also Lancet special credit - "Reprinted from The Lancet, Vol. number, Author(s), Title of article, Pages No., Copyright (Year), with permission from Elsevier."

4. Reproduction of this material is confined to the purpose and/or media for which permission is hereby given. The material may not be reproduced or used in any other way, including use in combination with an artificial intelligence tool (including to train an algorithm, test, process, analyse, generate output and/or develop any form of artificial intelligence tool), or to create any derivative work and/or service (including resulting from the use of artificial intelligence tools).

5. Altering/Modifying Material: Not Permitted. However figures and illustrations may be altered/adapted minimally to serve your work. Any other abbreviations, additions, deletions and/or any other alterations shall be made only with prior written authorization of Elsevier Ltd. (Please contact Elsevier's permissions helpdesk here). No modifications can be made to any Lancet figures/tables and they must be reproduced in full.

6. If the permission fee for the requested use of our material is waived in this instance, please be advised that your future requests for Elsevier materials may attract a fee.

7. Reservation of Rights: Publisher reserves all rights not specifically granted in the combination of (i) the license details provided by you and accepted in the course of this licensing transaction, (ii) these terms and conditions and (iii) CCC's Billing and Payment terms and conditions.

8. License Contingent Upon Payment: While you may exercise the rights licensed immediately upon issuance of the license at the end of the licensing process for the transaction, provided that you have disclosed complete and accurate details of your proposed use, no license is finally effective unless and until full payment is received from you (either by publisher or by CCC) as provided in CCC's Billing and Payment terms and conditions. If full payment is not received on a timely basis, then any license preliminarily granted shall be deemed automatically revoked and shall be void as if never granted. Further, in the event that you breach any of these terms and conditions or any of CCC's Billing and Payment terms and conditions, the license is automatically revoked and shall be void as if never granted. Use of materials as described in a revoked license, as well as any use of the materials beyond the scope of an unrevoked license, may constitute copyright infringement and publisher reserves the right to take any and all action to protect its copyright in the materials.

9. Warranties: Publisher makes no representations or warranties with respect to the licensed material.

10. Indemnity: You hereby indemnify and agree to hold harmless publisher and CCC, and their respective officers, directors, employees and agents, from and against any and all claims arising out of your use of the licensed material other than as specifically authorized pursuant to this license.

11. No Transfer of License: This license is personal to you and may not be sublicensed, assigned, or transferred by you to any other person without publisher's written permission.

12. No Amendment Except in Writing: This license may not be amended except in a writing signed by both parties (or, in the case of publisher, by CCC on publisher's behalf).

13. Objection to Contrary Terms: Publisher hereby objects to any terms contained in any purchase order, acknowledgment, check endorsement or other writing prepared by you, which terms are inconsistent with these terms and conditions or CCC's Billing and Payment terms and conditions. These terms and conditions, together with CCC's Billing and Payment terms and conditions (which are incorporated herein), comprise the entire agreement between you and publisher (and CCC) concerning this licensing transaction. In the event of any conflict between your obligations established by these terms and conditions and those established by CCC's Billing and Payment terms and conditions, these terms and conditions shall control.

14. Revocation: Elsevier or Copyright Clearance Center may deny the permissions described in this License at their sole discretion, for any reason or no reason, with a full refund payable to you. Notice of such denial will be made using the contact information provided by you. Failure to receive such notice will not alter or invalidate the denial. In no event will Elsevier or Copyright Clearance Center be responsible or liable for any costs, expenses or damage incurred by you as a result of a denial of your permission request, other than a refund of the amount(s) paid by you to Elsevier and/or Copyright Clearance Center for denied permissions.

LIMITED LICENSE

The following terms and conditions apply only to specific license types:

15. **Translation**: This permission is granted for non-exclusive world **English** rights only unless your license was granted for translation rights. If you licensed translation rights you may only translate this content into the languages you requested. A professional translator must perform all translations and reproduce the content word for word preserving the integrity of the article.

16. **Posting licensed content on any Website**: The following terms and conditions apply as follows: Licensing material from an Elsevier journal: All content posted to the web site must maintain the copyright information line on the bottom of each image; A hyper-text must be included to the Homepage of the journal from which you are licensing at http://www.sciencedirect.com/science/journal/xxxxx or the Elsevier homepage for books at http://www.elsevier.com; Central Storage: This license does not include permission for a scanned version of the material to be stored in a central repository such as that provided by Heron/XanEdu.

Licensing material from an Elsevier book: A hyper-text link must be included to the Elsevier homepage at http://www.elsevier.com . All content posted to the web site must maintain the copyright information line on the bottom of each image.

Posting licensed content on Electronic reserve: In addition to the above the following clauses are applicable: The web site must be password-protected and made available only to bona fide students registered on a relevant course. This permission is granted for 1 year only. You may obtain a new license for future website posting.

17. **For journal authors:** the following clauses are applicable in addition to the above:

Preprints:

A preprint is an author's own write-up of research results and analysis, it has not been peer-reviewed, nor has it had any other value added to it by a publisher (such as formatting, copyright, technical enhancement etc.).

Authors can share their preprints anywhere at any time. Preprints should not be added to or enhanced in any way in order to appear more like, or to substitute for, the final versions of articles however authors can update their preprints on arXiv or RePEc with their Accepted Author Manuscript (see below).

If accepted for publication, we encourage authors to link from the preprint to their formal publication via its DOI. Millions of researchers have access to the formal publications on ScienceDirect, and so links will help users to find, access, cite and use the best available version. Please note that Cell Press, The Lancet and some society-owned have different preprint policies. Information on these policies is available on the journal homepage.

Accepted Author Manuscripts: An accepted author manuscript is the manuscript of an article that has been accepted for publication and which typically includes author-incorporated changes suggested during submission, peer review and editor-author communications.

Authors can share their accepted author manuscript:

- immediately
    - via their non-commercial person homepage or blog
    - by updating a preprint in arXiv or RePEc with the accepted manuscript
    - via their research institute or institutional repository for internal institutional uses or as part of an invitation-only research collaboration work-group
    - directly by providing copies to their students or to research collaborators for their personal use
    - for private scholarly sharing as part of an invitation-only work group on commercial sites with which Elsevier has an agreement
- After the embargo period
    - via non-commercial hosting platforms such as their institutional repository
    - via commercial sites with which Elsevier has an agreement

In all cases accepted manuscripts should:

- link to the formal publication via its DOI
- bear a CC-BY-NC-ND license - this is easy to do
- if aggregated with other manuscripts, for example in a repository or other site, be shared in alignment with our hosting policy not be added to or enhanced in any way to appear more like, or to substitute for, the published journal article.

Published journal article (JPA): A published journal article (PJA) is the definitive final record of published research that appears or will appear in the journal and embodies all value-adding publishing activities including peer review co-ordination, copy-editing, formatting, (if relevant) pagination and online enrichment.

Policies for sharing publishing journal articles differ for subscription and gold open access articles:

Subscription Articles: If you are an author, please share a link to your article rather than the full-text. Millions of researchers have access to the formal publications on ScienceDirect, and so links will help your users to find, access, cite, and use the best available version.

Theses and dissertations which contain embedded PJAs as part of the formal submission can be posted publicly by the awarding institution with DOI links back to the formal publications on ScienceDirect.

If you are affiliated with a library that subscribes to ScienceDirect you have additional private sharing rights for others' research accessed under that agreement. This includes use

for classroom teaching and internal training at the institution (including use in course packs and courseware programs), and inclusion of the article for grant funding purposes.

Gold Open Access Articles: May be shared according to the author-selected end-user license and should contain a CrossMark logo, the end user license, and a DOI link to the formal publication on ScienceDirect.

Please refer to Elsevier's posting policy for further information.

18. **For book authors** the following clauses are applicable in addition to the above: Authors are permitted to place a brief summary of their work online only. You are not allowed to download and post the published electronic version of your chapter, nor may you scan the printed edition to create an electronic version. **Posting to a repository:** Authors are permitted to post a summary of their chapter only in their institution's repository.

19. **Thesis/Dissertation**: If your license is for use in a thesis/dissertation your thesis may be submitted to your institution in either print or electronic form. Should your thesis be published commercially, please reapply for permission. These requirements include permission for the Library and Archives of Canada to supply single copies, on demand, of the complete thesis and include permission for Proquest/UMI to supply single copies, on demand, of the complete thesis. Should your thesis be published commercially, please reapply for permission. Theses and dissertations which contain embedded PJAs as part of the formal submission can be posted publicly by the awarding institution with DOI links back to the formal publications on ScienceDirect.

Elsevier Open Access Terms and Conditions

You can publish open access with Elsevier in hundreds of open access journals or in nearly 2000 established subscription journals that support open access publishing. Permitted third party re-use of these open access articles is defined by the author's choice of Creative Commons user license. See our open access license policy for more information.

Terms & Conditions applicable to all Open Access articles published with Elsevier:

Any reuse of the article must not represent the author as endorsing the adaptation of the article nor should the article be modified in such a way as to damage the author's honour or reputation. If any changes have been made, such changes must be clearly indicated.

The author(s) must be appropriately credited and we ask that you include the end user license and a DOI link to the formal publication on ScienceDirect.

If any part of the material to be used (for example, figures) has appeared in our publication with credit or acknowledgement to another source it is the responsibility of the user to ensure their reuse complies with the terms and conditions determined by the rights holder.

Additional Terms & Conditions applicable to each Creative Commons user license:

CC BY: The CC-BY license allows users to copy, to create extracts, abstracts and new works from the Article, to alter and revise the Article and to make commercial use of the Article (including reuse and/or resale of the Article by commercial entities), provided the user gives appropriate credit (with a link to the formal publication through the relevant DOI), provides a link to the license, indicates if changes were made and the licensor is not represented as endorsing the use made of the work. The full details of the license are available at http://creativecommons.org/licenses/by/4.0.

CC BY NC SA: The CC BY-NC-SA license allows users to copy, to create extracts, abstracts and new works from the Article, to alter and revise the Article, provided this is not done for commercial purposes, and that the user gives appropriate credit (with a link to the formal publication through the relevant DOI), provides a link to the license, indicates if changes were made and the licensor is not represented as endorsing the use made of the work. Further, any new works must be made available on the same conditions. The full details of the license are available at http://creativecommons.org/licenses/by-nc-sa/4.0.

CC BY NC ND: The CC BY-NC-ND license allows users to copy and distribute the Article, provided this is not done for commercial purposes and further does not permit distribution of the Article if it is changed or edited in any way, and provided the user gives appropriate credit (with a link to the formal publication through the relevant DOI), provides a link to the license, and that the licensor is not represented as endorsing the use made of the work. The full details of the license are available at http://creativecommons.org/licenses/by-nc-nd/4.0. Any commercial reuse of Open Access articles published with a CC BY NC SA or CC BY NC ND license requires permission from Elsevier and will be subject to a fee.

Commercial reuse includes:

- Associating advertising with the full text of the Article
- Charging fees for document delivery or access
- Article aggregation
- Systematic distribution via e-mail lists or share buttons

Posting or linking by commercial companies for use by customers of those companies.

20. **Other Conditions**:

v1.10

Questions? customercare@copyright.com.

Publishing") consists of your order details and the terms and conditions provided by AIP Publishing and Copyright Clearance Center.

Order Number	501871652
Order date	Dec 30, 2023
Licensed Content Publisher	AIP Publishing
Licensed Content Publication	Physics of Fluids 1958-1988
Licensed Content Title	An electron conductivity model for dense plasmas
Licensed Content Author	Lee, Y. T.; More, R. M.
Licensed Content Date	May 1, 1984
Licensed Content Volume	27
Licensed Content Issue	5
Type of Use	Thesis/Dissertation
Requestor type	Student
Format	Print
Portion	Figure/Table
Number of figures/tables	1
Will you be translating?	No
Title of new work	Simulating radiatively cooled pulsed-power driven systems using static mesh refinement
Institution name	Imperial College London
Expected presentation date	Dec 2023
Portions	Fig 3
Requestor Location	Ms. Nikita Chaturvedi 28 Westminster Palace Gardens London, SW1P 1RR United Kingdom Attn: Ms. Nikita Chaturvedi
Billing Type	Invoice
Billing Address	Ms. Nikita Chaturvedi 28 Westminster Palace Gardens London, United Kingdom SW1P 1RR Attn: Ms. Nikita Chaturvedi
Total	0.00 USD

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

Print This Page