

Title	PION: simulating bow shocks and circumstellar nebulae
Creators	Mackey, Jonathan and Green, Samuel and Moutzouri, Maria and Haworth, Thomas J. and Kavanagh, Robert D. and Zargaryan, Davit and Celeste, Maggie
Date	2021
Citation	Mackey, Jonathan and Green, Samuel and Moutzouri, Maria and Haworth, Thomas J. and Kavanagh, Robert D. and Zargaryan, Davit and Celeste, Maggie (2021) PION: simulating bow shocks and circumstellar nebulae. Monthly Notices of the Royal Astronomical Society, 504. p. 983. ISSN 1365-2966
URL	https://dair.dias.ie/id/eprint/1124/
DOI	http://dx.doi.org/10.1093/mnras/stab781

PION: simulating bow shocks and circumstellar nebulae

Jonathan Mackey^{1,2}★, Samuel Green^{1,2}, Maria Moutzouri^{1,2}, Thomas J. Haworth³,
Robert D. Kavanagh⁴, Davit Zargaryan^{1,2,5} and Maggie Celeste^{1,4}

¹Dublin Institute for Advanced Studies, Astronomy and Astrophysics Section, 31 Fitzwilliam Place, Dublin 2, Ireland

²Dublin Institute for Advanced Studies, Centre for AstroParticle Physics and Astrophysics (CAPPA), DIAS Dunsink Observatory, Dunsink Lane, Dublin 15, Ireland

³Astronomy Unit, School of Physics and Astronomy, Queen Mary University of London, London E1 4NS, UK

⁴School of Physics, Trinity College Dublin, The University of Dublin, Dublin 2, Ireland

⁵High Energy Astrophysics Laboratory, RAU, 123 Hovsep Emin St, Yerevan 0051, Armenia

Accepted 2021 March 5. in original form 2021 January 8

ABSTRACT

Expanding nebulae are produced by mass-loss from stars, especially during late stages of evolution. Multidimensional simulation of these nebulae requires high resolution near the star and permits resolution that decreases with distance from the star, ideally with adaptive time-steps. We report the implementation and testing of static mesh-refinement in the radiation-magnetohydrodynamics (R-MHD) code PION, and document its performance for 2D and 3D calculations. The bow shock produced by a hot, magnetized, slowly rotating star as it moves through the magnetized ISM is simulated in 3D, highlighting differences compared with 2D calculations. Latitude-dependent, time-varying magnetized winds are modelled and compared with simulations of ring nebulae around blue supergiants from the literature. A 3D simulation of the expansion of a fast wind from a Wolf–Rayet star into the slow wind from a previous red supergiant phase of evolution is presented, with results compared with results in the literature and analytic theory. Finally, the wind–wind collision from a binary star system is modelled with 3D MHD, and the results compared with previous 2D hydrodynamic calculations. A PYTHON library is provided for reading and plotting simulation snapshots, and the generation of synthetic infrared emission maps using TORUS is also demonstrated. It is shown that state-of-the-art 3D MHD simulations of wind-driven nebulae can be performed using PION with reasonable computational resources. The source code and user documentation is made available for the community under a BSD3 licence.

Key words: hydrodynamics – radiative transfer – methods: numerical – stars: winds, outflows – ISM: bubbles.

1 INTRODUCTION

Massive stars emit copious extreme-ultraviolet (EUV) photons capable of ionizing hydrogen when on the hydrogen-burning main sequence and also have line-driven stellar winds with terminal velocities $v_\infty \gtrsim 1000 \text{ km s}^{-1}$ (Snow & Morton 1976), with important consequences for their surroundings (Dale 2015). After the main-sequence phase, the outer layers of a massive star expand rapidly, and the star evolves to the upper right-hand part of the Hertzsprung–Russell Diagram (HRD), becoming a cool and luminous supergiant. Such stars have extended, loosely bound envelopes, and their further evolution is determined by mass-loss through winds, eruptions, or interaction with a binary companion, and by rotational mixing of nuclear-processed material from the core to the envelope (Langer 2012; Smith 2014).

After the main sequence, the dynamical time-scale of circumstellar nebulae ($\sim 10^4$ – 10^5 yr) becomes comparable to the nuclear ($\sim 10^5$ yr) and thermal (~ 10 – 10^4 yr) time-scales of a massive star. Mass-loss rates (\dot{M}) and wind velocities (v_∞) can change drastically on these time-scales, meaning that the evolution of circumstellar

nebulae cannot be considered in isolation from the evolution of the central star(s). These late phases of evolution of massive stars are very uncertain because some key physical processes are poorly constrained and poorly modelled, namely convection, mass-loss, rotation, and interaction with a companion (for a review, see Smith 2014). Significant progress is being made in understanding the radii (Grassitelli et al. 2018) and wind structure (Sander & Vink 2020) of classical Wolf–Rayet (WR) stars, winds from stars close to the Eddington limit (Bestenlehner 2020) and potentially understanding the S-Doradus cycle of Luminous Blue Variables (LBVs) (Grassitelli et al. 2021). On the other hand, we do not yet have a predictive theory of mass-loss from red supergiants (RSGs), for which the empirical scaling of \dot{M} with stellar luminosity, mass, and temperature is a subject of active debate and research (Beasor et al. 2020; Humphreys et al. 2020). Nor is there any consensus on the causes or trigger for eruptive mass-loss events such as LBV giant eruptions, but there are indications that some of them could be driven by binary interaction (Smith et al. 2018). While we do not have a predictive theory for mass-loss rates across the HRD, stellar evolution calculations use mass-loss prescriptions that do cover the HRD, and so the wind-driven nebula produced around a massive star is a prediction of stellar evolution calculations. Comparing these predictions with observations is a test of mass-loss prescriptions.

* E-mail: jmackey@cp.dias.ie

Circumstellar nebulae are complex structures, typically subject to non-linear dynamical instabilities (García-Segura, Langer & Mac Low 1996b), and must be studied with multidimensional radiation-magnetohydrodynamics (R-MHD), or simplifications thereof (e.g. hydrodynamics (HD) or ideal magnetohydrodynamics (MHD) with radiative heating and cooling). This means that, while significant work has been done on modelling circumstellar nebulae (see below), its potential to test stellar evolution theory has not been exploited to the extent that it could be.

The first two-dimensional (2D) hydrodynamics simulations of the expansion of H II regions (Bodenheimer, Tenorio-Tagle & Yorke 1979), wind bubbles (Rozyczka 1985), superbubbles (Mac Low, McCray & Norman 1989) and bow shocks (Mac Low et al. 1991) showed the importance of asymmetric ISM density and of hydrodynamical instabilities in the evolution of circumstellar nebulae. Colliding winds in binary systems were studied by Stevens, Blondin & Pollock (1992), who found that the wind collision region can be dynamically unstable and predicted that the resulting X-ray emission could vary at the level of 10 per cent. A series of papers gave a quantitative understanding of the physical processes that could give rise to the structure of planetary nebulae (Frank & Mellema 1994; Mellema 1994). Raga et al. (1997) studied the properties of bow shocks and H II regions around runaway stars.

By coupling stellar-evolution calculations with 2D hydrodynamics on a spherical mesh (with logarithmic radial spacing), García-Segura et al. (1996a,b) studied the development of nebulae around stars evolving from main sequence through LBV \rightarrow WR, and RSG \rightarrow WR phases, respectively. They predicted lifetimes and observable properties of nebulae produced during various transitions and phases, and compared results with nebulae around a number of WR stars, finding good agreement in some cases. The numerical methods developed have been used in many follow-up works and ported to other codes (e.g. van Marle, Langer & García-Segura 2005; Chita et al. 2008; van Veelen et al. 2009; van Marle, Decin & Meliani 2014).

Meyer et al. (2014) implemented the wind boundary-condition and radiative heating/cooling model of Mackey et al. (2012) into PLUTO (Mignone et al. 2012) and made 2D hydrodynamics simulations of bow shocks around massive stars moving through the Galactic plane, following this in Meyer et al. (2015) with simulations of supernova blastwaves interacting with the bow shocks. This model was extended to MHD by Meyer et al. (2017) and also used for a number of recent studies of circumstellar nebulae (e.g. Meyer, Petrov & Pohl 2020).

Yorke & Kaisig (1995) and Yorke & Welz (1996) developed a radiation-hydrodynamics (R-HD) solver on a multiply nested grid in 2D cylindrical coordinates (R and z) with adaptive time-steps. This was used by Freyer, Hensler & Yorke (2003, 2006) to study H II regions and wind bubbles around two stars (60 and 35 M_{\odot} , respectively) for the full evolution of the star through main sequence, supergiant, and WR phases. The same evolutionary tracks as García-Segura et al. (1996a,b) were used.

Three-dimensional (3D) simulations of circumstellar nebulae became possible in the past 10–15 yr. Pittard (2009) developed 3D hydrodynamic simulations including wind acceleration, used to study the thermal X-ray emission from binary stars in Pittard & Parkin (2010). Using 3D adaptive mesh-refinement simulations, Parkin et al. (2011) studied the wind-collision region of the binary system η Carinae, and Parkin & Gosset (2011) investigated the WR 22 system. H II region expansion in turbulent clouds was investigated by a number of authors (Mellema et al. 2006b; Arthur et al. 2011; Walch et al. 2012; Geen et al. 2015b). 3D simulations of bow shocks around RSGs were presented in Mohamed, Mackey & Langer

(2012). The differences between 2D and 3D calculations of wind–wind interaction were investigated by van Marle & Keppens (2012), and 3D calculations of wind bubbles expanding in turbulent media by Rogers & Pittard (2013). Geen et al. (2015a) and Haid et al. (2018) studied combined effects of winds and H II regions on the ISM for the full evolution of a star using R-HD, similar to previous 2D calculations by Freyer et al. (2003). 3D MHD calculations of wind bubbles were presented by Scherer et al. (2020).

Previous work (Freyer et al. 2003, 2006) has shown the value of static mesh-refinement for simulating circumstellar structures expanding from small to large scales, motivating the work presented here. The majority of the work cited above was performed using software that is no longer actively developed or is not freely available. In this paper, we describe the simulation code *PhotoIonization of Nebulae*, abbreviated to PION, an R-MHD code that has been developed with the aim of modelling nebulae around massive evolving stars. Significant new additions to the code with respect to previous versions (Mackey & Lim 2010, 2011; Mackey 2012) are described, and the code is made available to the community under a BSD-3 licence from <https://www.pion.ie>.

The paper is organized as follows: Section 2 describes the numerical methods, including the wind boundary condition, static mesh-refinement, radiative transfer, and MHD. Test calculations are presented in Section 3 that show the strengths and weaknesses of static mesh-refinement. Applications of the code to modelling circumstellar nebulae are presented in Section 4, namely a 3D simulation of a magnetized bow shock (Section 4.1), 2D simulation of the formation of ring nebulae around rotating and evolving stars (Section 4.2), 3D R-HD simulation of the wind–wind interaction from an RSG evolving to a WR star (Section 4.3), and 3D MHD simulation of a wind–wind collision between two rotating stars (Section 4.4). In all of these cases, the results are compared with previous calculations in the literature. Methods for post-processing simulation snapshots are described in Section 5, and parallel scaling in Section 6. Conclusions are presented in Section 7.

2 CODE DESCRIPTION AND ALGORITHMS

PION is a HD and MHD grid-based simulation code that includes radiative transfer of ionizing and non-ionizing photons for R-HD (Mackey & Lim 2010) and for R-MHD (Mackey & Lim 2011). A finite-volume integration scheme was implemented that is second-order-accurate in time and space, following Falle, Komissarov & Joarder (1998). In Mackey & Lim (2010), Mackey & Lim (2011) the formation of pillars at the boundaries of H II regions was investigated using 3D simulations in Cartesian geometry. Improvements to the radiative transfer and time-integration schemes were described in Mackey (2012). 2D simulations with axisymmetry (cylindrical coordinates in R and z) were added following Falle (1991), and a stellar wind boundary condition implemented and used in Mackey et al. (2012) to study the nebula around Betelgeuse assuming the star was previously a blue supergiant and only recently evolved to a RSG. This was achieved by varying the wind parameters according to results from a stellar evolution calculation. The spherically symmetric (1D) coordinate system has also been implemented, and was used for studying the external irradiation of winds from RSG (Mackey et al. 2014; Szécsi, Mackey & Langer 2018) and for modelling the D-type expansion of H II regions (Bisbas et al. 2015). A non-equilibrium-ionization model for the thermodynamics and ionization of the diffuse ISM was introduced in Mackey, Langer & Gvaramadze (2013) for modelling H II regions, and a related model for molecular gas in Mackey et al. (2015), based on results from Henney et al.

(2009). These were used in Mackey et al. (2016), Gvaramadze et al. (2017), and Green et al. (2019) for simulating circumstellar nebulae and comparing observational data with synthetic observations.

These calculations were run on a uniform rectilinear grid, decomposed into blocks for parallel code execution, using MPI for inter-process communication. The code was shown to scale well to at least 256 cores for 2D problems, and to 1024 cores for 3D problems (Mackey 2012) in tests of strong scaling (i.e. fixed problem size, variable number of MPI processes). PION has proven to be a useful code for studying circumstellar nebulae and expanding bubbles driven by photoionization and winds, but most applications have been 2D because of the limitations of the uniform computational grid.

2.1 Stellar-wind boundary condition

A stellar wind is modelled as a source of mass, momentum, and energy within a sphere of user-specified radius on 1D-spherical, 2D-cylindrical, and 3D-Cartesian grids. There are three options in PION specified by wind-type parameter 0, 1, or 2 in the input parameter-file. Type 0 is a spherically symmetric wind that is constant in time; type 1 is a spherically symmetric and time-varying wind with properties specified by a text file containing the time evolution; and type 2 is a latitude-dependent and time-varying wind. These are described in the following subsections and demonstrated in Sections 4.1, 4.3, and 4.2, respectively. There is no limit on the number of wind sources that can be included in a simulation.

For all wind-boundary types, it is possible to specify chemical element abundance fractions (by mass) as passive scalar variables that are advected across the simulation domain (see also e.g. Georgy et al. 2013). These can be constant in time or with time-varying values read in from a text file. We implemented the consistent multispecies advection (sCMA) scheme of Plewa & Müller (1999) for tracking the fractional abundances of these chemical elements. This ensures that the non-uniform elemental abundances are tracked accurately as they expand outwards and mix with fluid elements that have (potentially) different abundances.

2.1.1 Constant wind

The simplest wind boundary (type 0) is spherically symmetric and constant in time, and the wind is injected at the terminal velocity, i.e. it is assumed that the wind boundary region is significantly bigger than the star. The boundary region is specified by a position and a radius, both in cm, and the physical properties of the wind are specified by the mass-loss rate, \dot{M} , the wind terminal velocity, v_∞ , stellar radius, R_* , temperature, T_{eff} , equatorial rotation velocity, v_{rot} , surface split-monopole magnetic field strength, B_* , and the mass fractions of any chemical elements tracked. Generally, the wind boundary region should be 10–20 grid cells in radius to suppress grid-related artefacts in the expanding flow.

If $|v_{\text{rot}}| > 0$ then the spherical symmetry is broken for multidimensional simulations, because the azimuthal component of velocity and magnetic field are non-zero. The magnetic field is taken to be weak (dynamically) and to follow a split monopole swept into a Parker spiral at large distance from the stellar surface. Both toroidal and poloidal field components are included, and for simplicity, it is assumed that the rotation and magnetic axes are coincident. The rotational component of the wind velocity decays with distance, r , from the star as r^{-1} and is generally negligible. The boundary condition follows closely the methods commonly used for MHD

modelling of the Solar Wind and Heliosphere (e.g. Pogorelov, Zank & Ogino 2004), also similar to the recent implementation on a spherical coordinate grid by Scherer et al. (2020), and it is demonstrated in Sections 4.1 and 4.4.

2.1.2 Time-varying wind

Wind type 1 is an extension of type 0 for time-varying sources that are specified through a tab-separated text file containing the evolution of the star in question. The columns in this file are: time, mass, luminosity, temperature, mass-loss rate, rotation velocity, critical rotation velocity, $v_{\text{crit}} \equiv v_{\text{esc}}/\sqrt{2}$ (where v_{esc} is the surface escape velocity), wind terminal velocity, and mass-fractions of any chemical elements tracked. All values are assumed to be in cgs units and can be modified output from a stellar evolution calculation (e.g. Mackey et al. 2012) or an ad-hoc model (cf. Langer, García-Segura & Mac Low 1999). The evolving stellar wind module was previously used in Mackey et al. (2012) to study the hydrodynamics of the nebula produced when a blue supergiant evolves redward to an RSG, and follows similar algorithms from the earlier literature (García-Segura et al. 1996b; van Marle et al. 2005). Here, the module is demonstrated in Section 4.3 for the nebula produced when an RSG evolves to a WR star.

2.1.3 Latitude-dependent and time-varying wind

Wind type 2 provides a prescription for latitude-dependent winds from rotating stars, and the option to read time-evolution of stellar-wind and radiation properties from a text file. The latitude-dependent wind is modelled following Langer et al. (1999), who introduced a mathematical model of the focusing of stellar wind towards the equator as the star approaches the so-called Ω -limit (Langer 1997), defined as the equatorial surface rotation speed, v_{rot} , for which the net acceleration on the surface layers is zero. The critical rotation velocity is used to define the rotation parameter, $\Omega \equiv v_{\text{rot}}/v_{\text{crit}} < 1$. Equations (3)–(5) in Langer et al. (1999) are used to calculate the latitude dependence of the wind density and velocity as a function of Ω . This algorithm is based on the theory of Bjorkman & Cassinelli (1993) and it produces many of the observed features of bipolar nebulae (Langer et al. 1999; Chita et al. 2008; van Marle et al. 2008), particularly for stars that reach critical rotation in the temperature range of 6000 – 10 000 K when embarking on a blue loop. This module is demonstrated in Section 4.2, where the results are compared with previous literature results.

2.2 Upgraded magnetohydrodynamics implementation

PION has an MHD implementation presented in Mackey & Lim (2011), which is effective for simulating the magnetohydrodynamics of H II regions (Mackey et al. 2013). This uses a modified version of the Dedner et al. (2002) mixed-GLM divergence-cleaning algorithm for mitigating against the growth of magnetic monopoles. It uses either the linear MHD solver described by Falle et al. (1998), or the Roe solver in conserved variables of Cargo & Gallice (1997), following Stone et al. (2008).

Neither of these MHD Riemann solvers is robust enough for the high-Mach-number shocks encountered in stellar-wind bubbles around hot stars. We implemented the HLL solver in HD (Harten, Lax & Leer 1983) and MHD (Janhunen 2000) following Mignone et al. (2012), and also the more accurate HLLD solver (Miyoshi &

Kusano 2005) for MHD. The HLLD solver is also not sufficiently robust for the high-Mach-number flows in stellar wind simulations, because it is not positive definite in gas pressure (Mignone et al. 2012). Following Mignone et al. (2012), we implemented a shock detection scheme and a switch that locally replaces the HLLD with the HLL scheme, which is positive definite. This improves the code stability, but for some problems the simpler HLL scheme should be used everywhere.

Following Derigs et al. (2018), we included the Powell source terms (Powell et al. 1999), added the source terms for the Dedner et al. (2002) ψ field, which we re-scaled as in Derigs et al. (2018), and included ψ in the total energy density. This introduces very small changes in the solution to test calculations, and some improvements in the robustness of the scheme. We did not implement the full scheme of Derigs et al. (2018) with their entropy-stable Riemann solver, and so the rest of the MHD implementation is as in Mackey & Lim (2011).

2.3 Static mesh-refinement

A number of implementations of static mesh-refinement have been described in the literature, typically arranged as a multiply nested grid that is centred on a region of interest. Freyer et al. (2003, 2006) used a 2D nested grid with axisymmetry described in Yorke & Kaisig (1995) and Yorke & Welz (1996) to study expanding nebulae. This has one advantage over a spherical grid with logarithmically spaced radial cells, in that the latter has a global time-step for all cells whereas the nested grid can have adaptive time-stepping. This efficiency comes at the cost that all radial columns away from a point source are not equal – angle-dependent numerical viscosity and grid-artefacts are inevitably introduced, as can be seen by comparing results from García-Segura et al. (1996b) and Freyer et al. (2006). In particular with 2D simulations, the symmetry axis has a coordinate singularity that affects results, also seen in bow-shock simulations (Green et al. 2019). For 3D, this is less of a problem, but the viscosity of the numerical scheme for expansion along grid axes remains different from expansion at an angle to the grid.

Recently Stone et al. (2020) described the implementation of static and adaptive mesh-refinement algorithms in the ATHENA++ software framework, again demonstrating the dramatic improvements that can be obtained with these techniques. The advantages in computational efficiency of a nested grid compared with a uniform grid are clear: for 2D calculations with a uniform grid, doubling the resolution everywhere increases the computational cost by a factor of 8; for 3D calculations, it is a factor of 16. Adding a nested grid that is a factor of 2 smaller than the coarse grid in each dimension, but that retains the same number of zones, increases the computational load by a factor of 3 (the fine grid requires the same amount of computation as the coarse grid *per step*, but must take twice the number of steps), and this is independent of dimensionality. Adding a third level requires seven times more computation than just a single level, whereas for a uniform grid the cost of quadrupling the resolution would be $64\times$ (2D) or $256\times$ (3D) more work. A nested grid also has a modest efficiency advantage over spherical-coordinate grids with a cell size that increases with radius, in that adaptive time-steps can be used. On a spherical grid all cells must use the same time-step, usually dictated by the smallest cells close to the origin. For N refinement levels, the computational saving using adaptive time-stepping compared with a global time-step on all levels approaches a factor of $N/2$ for large N .

There are three additions to a uniform-grid algorithm required for a nested grid:

(i) The refined grid should obtain its external boundary data from its parent (coarser) grid, by interpolating the coarse-grid zones to the zone-centres of the boundary data on the refined grid. This interpolation should be done to the same order of accuracy as the spatial reconstruction used, and should conserve the total mass, momentum and energy of the coarse-grid zone. This is known as *prolongation* (e.g. Tóth & Roe 2002).

(ii) The coarse grid should update its zones by obtaining averaged data from any finer-level grid (where applicable). This is known as *restriction* (e.g. Tóth & Roe 2002).

(iii) The flux entering/leaving a finer-level grid should be recorded and sent to the coarser-level parent grid to ensure that this flux is consistent across all grid levels (Berger & Colella 1989, hereafter BC89). This is required so that conserved quantities are indeed conserved; otherwise mass, momentum, and energy can disappear because of inconsistencies between levels.

All of these are well-established techniques, but they are described below because the implementation depends on the time-integration scheme adopted as well as the parallelization strategy.

2.3.1 Coarse-to-fine interpolation (Prolongation)

We follow the scheme used for MPI-AMRVAC (Meliani et al. 2007; Keppens et al. 2021) on a cell-by-cell basis, and for a grid with D spatial dimensions. For a scheme that is first-order accurate in space, we can simply copy the coarse-grid values to the 2^D fine-grid cells. For a second-order scheme, linear interpolation and correction are applied as follows:

(i) For each coarse-grid cell, i , with cell volume V_i , and cell-centred vector of primitive variables \mathbf{P}_i , calculate slopes, \mathbf{m}_k , of the primitive variables in each dimension k .

(ii) Send these data to the finer grid and the finer grid receives the data.

(iii) Using the slopes \mathbf{m}_k , interpolate \mathbf{P}_i to the cell centres of the 2^D fine-grid cells contained within the coarse-grid cell i , assigning primitive variable data \mathbf{P}_j to these cells. Depending on grid dimensionality, this uses linear, bilinear, or trilinear interpolation.

(iv) The conserved quantities \mathbf{U}_i and \mathbf{U}_j are calculated from \mathbf{P}_i and \mathbf{P}_j , respectively.

(v) The difference vector $\Delta = \mathbf{U}_i V_i - \sum_j \mathbf{U}_j V_j$ is calculated, to ensure that the conserved quantities have consistent values within the same volume in both levels.

(vi) The fine-grid cells j are corrected by adding $1/2^D$ of this difference to each \mathbf{U}_j (also dividing out the total volume)

$$\mathbf{U}_j \rightarrow \mathbf{U}_j + \frac{\Delta}{2^D V_i}. \quad (1)$$

(vii) The fine-grid primitive vectors \mathbf{P}_j are obtained from the corrected \mathbf{U}_j vectors, for each fine-grid cell.

This ensures that conserved quantities are conserved when a coarse grid cell is prolonged on to the finer grid.

2.3.2 Fine-to-coarse averaging (restriction)

This is much more straightforward than prolongation, and also independent of the spatial order of accuracy of the scheme.

(i) For each set of 2^D fine-grid cells j , contained within the coarse-grid cell i , we calculate the average of the conserved quantities

contained within the volume V_i of cell i

$$\mathbf{U}_i = \frac{\sum_j \mathbf{U}_j V_j}{\sum_j V_j}. \quad (2)$$

(ii) The list of \mathbf{U}_i vectors is sent to the coarse grid and the coarse grid receives data.

(iii) Vectors \mathbf{U}_i are converted to a primitive vector and assigned to each coarse-grid cell i .

2.3.3 Flux correction on coarse grid zones abutting a fine grid boundary

BC89 describe a method to ensure consistent fluxes across cell boundaries at different levels of refinement, with the assumption that the most accurately calculated flux is at the finest level. This finest-level flux is then propagated to coarser levels as needed, and the coarse-cell fluxes are corrected to agree with the finest-level flux. The PION implementation is described here for two levels, which is the only case that arises for a nested grid arrangement. It is assumed that the coarse and fine grids are assigned to different MPI processes, although the update algorithms do not make the MPI calls if the grids are on the same process.

The correction is not needed for the half-step in the second-order scheme, because this is only an approximate time-centred state used to calculate fluxes that are accurate to second order. This means that the full-step fluxes over two fine-grid steps must be added together and sent to the coarse grid *after* the full-step coarse fluxes have been calculated but *before* the coarse grid cells have their state advanced in time.

When the coarse and fine grids are set up, the edge cells of the refined levels are identified, as well as the coarse grid cells that share the same edge. Any fine-grid cells whose outer face is the edge of the fine level have up to D (for edge/corner cells) extra state vectors allocated and initialized to zero, to record the full-step fluxes as they are calculated.

In addition, the fine grid allocates up to 2D vectors of C-style *structs* (one for each outward normal direction on the grid). Each element in each vector represents a cell face on the coarse grid for which the flux will be corrected by the fine grid. The structs contain a list of pointers to the fine-grid cells contributing to this coarse-cell face, a vector of corresponding areas of the cell faces (not all identical for curvilinear grids), and a state vector to hold the time-integrated flux through the fine-cell faces over the two time steps. The coarse grid allocates a similar vector of structures to record the uncorrected fluxes, but there is only one coarse cell and one face area in each struct.

The scheme is as follows, shown only for one coordinate direction

(i) At the start of an even-numbered fine grid time-step, reset BC89 fluxes to zero.

(ii) Record fluxes, \mathbf{F}_j across all fine-grid boundary cells j during the time-step.

(iii) Calculate time- and area-integrated flux, $\Delta \mathbf{U}_i^f$, through surfaces of the 2^{D-1} cells, j , on the fine grid that correspond to the surface of coarse-grid cell, i . Add these to the vector of structs on fine grid by summing the fluxes

$$\Delta \mathbf{U}_i^f = \Delta t^f \sum_{j=0}^{2^{(D-1)}} \mathbf{F}_j A_j^f, \quad (3)$$

where A_j^f is the surface area of the face of cell j , and Δt^f is the fine-grid time-step.

(iv) On the coarse grid record fluxes, \mathbf{F}_i , through cell faces that map on to the edge of fine grid, and calculate $\Delta \mathbf{U}_i^c = \Delta t^c \mathbf{F}_i A_i^c$ (where $\Delta t^c = 2\Delta t^f$ is the coarse-grid time step and A_i^c is the area of the face of cell i).

(v) Complete the odd-numbered fine-grid step by repeating steps (ii) and (iii), adding to $\Delta \mathbf{U}_i^f$ as before.

(vi) Send array of $\Delta \mathbf{U}_i^f$ values from MPI process of fine grid to MPI process of coarse grid.

(vii) Correct $\Delta \mathbf{U}_i^c$ values on the coarse grid so they agree with $\Delta \mathbf{U}_i^f$, and modify fluxes accordingly.

2.4 Time-integration scheme

The coarse-to-fine update can only be applied once every full step of the coarse grid for a fully consistent solution, and so the finer-level grid must calculate two time-steps between updates, following the algorithm above. The boundary region should be six cells deep in order to complete two full steps on the finer level without updating the boundary conditions (for a second-order scheme), compared with four cells if the update was every step. The fine-to-coarse boundaries are updated every full step on the fine-level grid.

The following time-integration scheme was implemented, based on the uniform-grid scheme of Falle et al. (1998) and using adaptive time-steps on nested grids. We update level l by one step and level $l + 1$ by two steps, and the algorithm is recursive.

- (i) Begin time-step level l , to advance current time, t_0 , by Δt_l .
- (ii) If an even step, receive coarse-to-fine external boundary data from level $l - 1$.
- (iii) Update any other external boundary conditions (including from domain decomposition).
- (iv) Send coarse-to-fine data to $l + 1$.
- (v) Advance level $l + 1$ by one step.
- (vi) Calculate fluxes on level l and calculate the time-centred state at $t_0 + 0.5\Delta t_l$.
- (vii) Update internal boundary conditions (e.g. stellar wind properties).
- (viii) Receive fine-level data from $l + 1$ and replace level l states with these data (including optical depths, if raytracing).
- (ix) Update external boundary conditions except for coarse-to-fine level boundary.
- (x) Do raytracing on time-centred state to calculate optical depths for full step.
- (xi) Calculate level l fluxes for full step, using the time-centred state, saving fluxes needed for BC89 correction.
- (xii) Advance level $l + 1$ by one step.
- (xiii) Receive BC89 fluxes summed over two steps from $l + 1$ and correct level l fluxes accordingly.
- (xiv) Update state vector on level l to $t = t_0 + \Delta t_l$.
- (xv) Update internal boundary conditions (e.g. stellar wind properties), receive fine-level data from $l + 1$ (including optical depths) and replace level l states with these data.
- (xvi) Raytrace level l to calculate optical depths for next (half) step.
- (xvii) If an even-numbered step, send BC89 fluxes to $l - 1$.
- (xviii) Send fine-to-coarse averaged data to $l - 1$ (including optical depths).
- (xix) Return.

After a full step on the coarsest grid, the time-step is re-calculated on all grids and a new step is started. The refined grids use the same

time-step for the duration of the coarse step, and so a safety factor is included in addition to the usual CFL condition.

2.5 Radiative Transfer

Raytracing is implemented in PION using the method of short characteristics with the On-The-Spot approximation, i.e. scattered radiation is locally re-absorbed and so only direct radiation from point sources needs to be transported across the grid (Mellema et al. 2006a; Mackey 2012). If radiation sources are always on the most refined grid, then raytracing on a nested grid proceeds on the finest level exactly as for a uniform grid. On the next coarser level, the stored quantities (whether optical depth, column density, or radiation density) are mapped on to the coarser grid cells from the finer grid, and raytracing proceeds through the rest of the level. This procedure is repeated for all coarse levels.

Raytracing must be performed on all levels when calculating the time-step on the coarsest level (because the photoionization and recombination time-scales also limit the time-step), and then twice each time-step per level for the second-order scheme (Mackey 2012). So for a grid with four levels, the finest level (level 3) has 17 raytracings per coarse-grid step, level 2 has 9 raytracings, level 1 has 5, and level 0 has 3.

2.6 Summary

The upgraded MHD algorithms with static mesh-refinement are implemented in first-order and second-order schemes, and the code was run on different numbers of MPI processes to check for consistency. The results for the first-order scheme for HD, MHD, and R-MHD were shown to be byte-for-byte identical, independent of the number of MPI processes. For the second-order scheme, HD and MHD results are byte-for-byte identical when run on different numbers of MPI processes, and R-MHD results show very small differences (relative difference $\approx 10^{-4}$ in primitive variables) at the end of a simulation, arising because of the adaptive time-stepping algorithm in the implicit solver for ionization and heating/cooling. The conservation of mass, momentum, and energy were also checked and found to be consistent with roundoff error.

3 TEST CALCULATIONS

3.1 Advection of a magnetic field loop

Advection across refinement boundaries can verify that the refinement has been implemented correctly, and can show that the accuracy of the nested-grid integration algorithm is the same as that of the uniform-grid. This is demonstrated with 2D test problems using periodic boundaries, where the whole domain is advected twice through the domain and back to its starting location.

The advection of a magnetic field loop is a good test of the diffusivity of an MHD scheme (e.g. Gardiner & Stone 2005; Stone et al. 2008). A weak magnetic field loop is set up in the x - y plane using the vector potential $\mathbf{A} = [0, 0, A_z]$, with z -component

$$A_z = \begin{cases} \sqrt{4\pi}A_0(R_0 - r) & (r = \sqrt{x^2 + y^2}) < R_0 \\ 0 & r \geq R_0, \end{cases} \quad (4)$$

using $A_0 = 0.001$ and $R_0 = 0.3$. This generates a constant circular magnetic field of strength $\sqrt{4\pi}A_0$ within $r < R_0$, a current sheet at $r = R_0$, and a current spike at $r = 0$ whose amplitude increases with increasing numerical resolution. The initial uniform density is $\rho = 1$, thermal pressure is $p_g = 1$, and velocity is $\mathbf{v} = [2, 1, 0]$, using

an adiabatic equation of state with $\gamma = 5/3$. The magnetic pressure, $p_m \equiv B^2/8\pi = 5 \times 10^{-7}$, is therefore negligible and the field is advected with the flow.

For the uniform-grid simulation, a 2D domain with $x \in [-1, 1]$ and $y \in [-0.5, 0.5]$ is used, and for the nested-grid simulation the domain is $x \in [-1, 1]$ and $y \in [-1, 1]$ (so that the field loop fits entirely in the first refined level centred on $[0,0]$). In both cases, the HLLD solver is used, and 200×100 grid cells per level. Results from the uniform-grid simulation are plotted in Fig. 1, showing the initial conditions (left-hand panel) and the final state at $t = 2$ (right-hand panel), by which time the loop has advected twice across the domain at an angle 30° to the positive x -axis. These results are similar to those shown for the previous version of PION in Mackey & Lim (2011) using the linear solver of Falle et al. (1998). The decay of magnetic pressure very closely follows the results presented in Mackey & Lim (2011) (because the integration scheme is essentially unchanged) and is not shown here.

Results for the simulation with one level of refinement, and a refined grid on $x \in [-0.5, 0.5]$ and $y \in [-0.5, 0.5]$, are plotted in Fig. 2, where now contours are only plotted for the refined grid. This simulation has a larger domain from $[-1, 1]$ to $[1, 1]$. The loop still advects twice across the domain, crossing both the refined and coarse levels, but spending most time on the coarse level. Results are very similar to the uniform-grid case, except that the extrema of $\nabla \times \mathbf{B}$ are slightly more pronounced, meaning that the initial conditions are marginally better preserved. No artefacts are introduced by advecting the field loop across refinement levels.

3.2 MHD blast wave in 2D

The expansion of a blastwave in a 2D Cartesian domain is a standard test problem, (e.g. Stone et al. 2008). Here, we set up the problem as in Stone et al. (2008) and Mackey & Lim (2011): the domain is $x \in [-0.5, 0.5]$, $y \in [-0.75, 0.75]$, resolved by 256×384 cells, with a uniform background density $\rho = 1$, pressure $p_g = 0.1$, and magnetic field strength of 0.1, 1.0, or 10.0 (in units where factors of 4π do not appear, so e.g. $B = 1$ corresponds to $B = \sqrt{4\pi}$ in cgs units). The field is oriented at an angle of 45° to the x -axis and the medium is initially at rest. A circle of radius 0.1 is filled with gas at pressure $p_g = 10$ and the system is allowed to evolve to $t = 0.2$. Periodic boundary conditions are imposed on all sides, although they are not relevant to the dynamics until $t > 0.2$.

We calculate the three problems (weak, medium, and strong magnetic field) using the HLLD solver with a CFL number of 0.24, initially on a uniform grid. The results for all three cases are shown in Fig. 3 at $t = 0.2$, and are comparable to results obtained with ATHENA (Stone et al. 2008). The features visible in the contour at $\rho = 1$, outside the outer shock, arise from the diffusion of divergence errors by the ψ field of the GLM-MHD scheme. Apart from this, the symmetry of the blast wave and contact discontinuity are maintained well, and the HLLD solver is at least as good the solver presented in Mackey & Lim (2011) while being significantly more robust.

For comparison, the results using a nested grid with two levels, centred on the origin, for all three magnetic-field strengths are plotted in Fig. 4, using the same colour scale and contours as Fig. 3. The inner part of the solution is solved more accurately, as expected because of the higher resolution, but the most obvious difference from Fig. 3 is that the refinement boundary has left an imprint in the form of waves trailing the forward shock in an approximate parallelogram shape in the middle panel ($B = 1$). For the case where $B = 0.1$, this grid-refinement-boundary effect is much less noticeable than for $B = 1$. The effect is almost absent in HD calculations, similar to panel

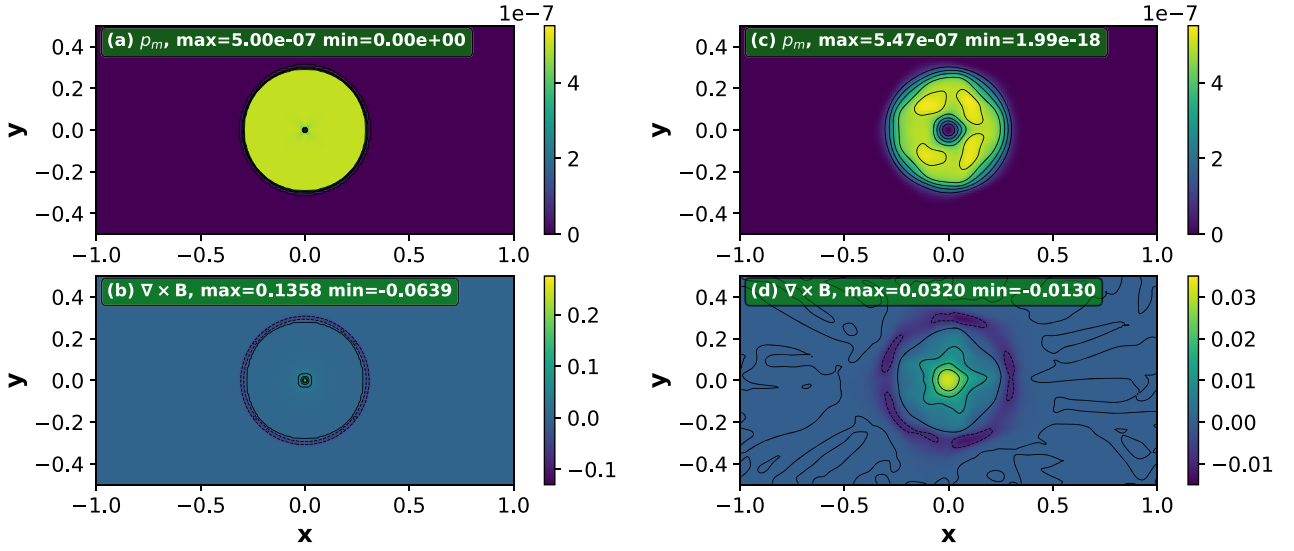


Figure 1. Field-loop advection test, showing the initial conditions (panels a,b) and the final state at $t = 2$ after advecting twice across the domain (panels c,d). The magnetic pressure $p_m = |\mathbf{B}|^2/(8\pi)$ is plotted in the upper panels, and the current density in the lower panels, using a linear colour scale as indicated. In the upper panels, contours of magnetic pressure are shown from $p_m = 1 \times 10^{-7}$ to 5×10^{-7} , linearly spaced in steps of 1×10^{-7} . For panel (b), the current density contours are $\nabla \times \mathbf{B} = [-0.06, -0.03, 0, 0.03, 0.06, 0.09, 0.12]$, and for panel (d) $\nabla \times \mathbf{B} = [-0.008, 0, 0.008, 0.016, 0.024, 0.032]$, using broken lines for negative contours.

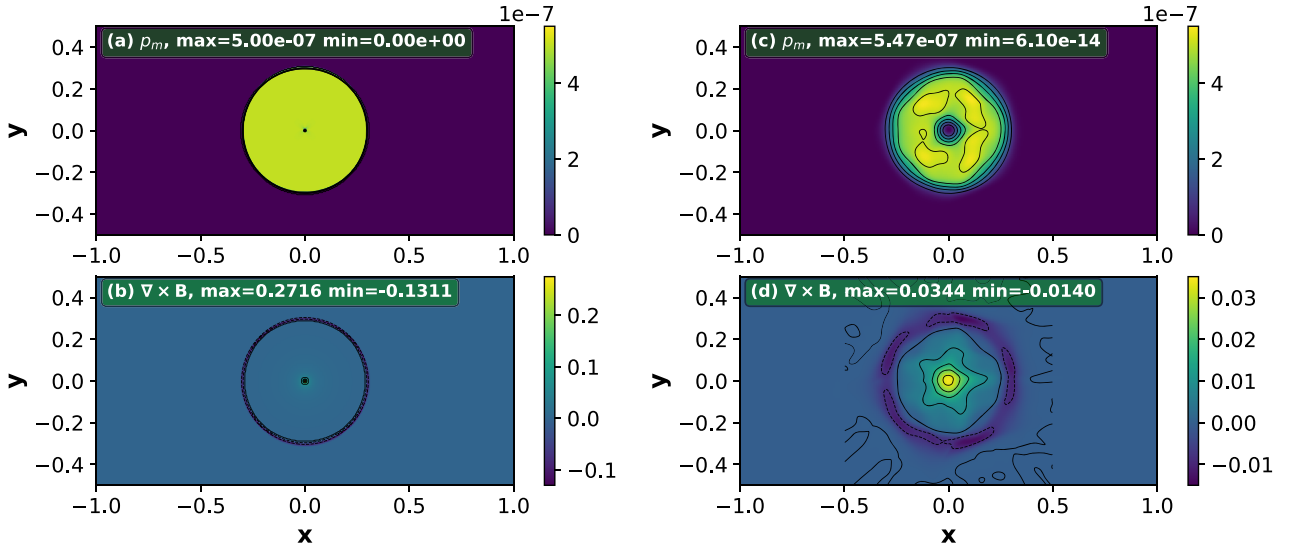


Figure 2. As Fig. 1, but using a nested grid with two levels, centred on $[0,0]$, again plotting magnetic pressure above and current density below. The left-hand panels show the initial conditions and the right-hand panels the results at $t = 2$ after advecting twice across the domain. In the upper panels, contours of magnetic pressure are shown as before, but only for the refined grid. For panel (b), the current density contours are $\nabla \times \mathbf{B} = [-0.12, -0.06, 0, 0.06, 0.12, 0.18, 0.24]$, and for panel (d) $\nabla \times \mathbf{B} = [-0.008, 0, 0.008, 0.016, 0.024, 0.032]$, using broken lines for negative contours.

(a) in Fig. 4 for which p_g in the hot region is much larger than the magnetic pressure, $p_m \equiv B^2/2$ (in these units), i.e. the plasma $\beta \equiv p_g/p_m \gg 1$. In panel (c) of Fig. 4, the grid effect is visible in that the blastwave is no longer mirror-symmetric along its axis, but the error is not worse than in panel (b). In panel (b), the initial conditions have $\beta \ll 1$ in the undisturbed medium, and $\beta = 2$ in the hot region of the initial conditions. Panel (c) has $\beta < 1$ in all regions of the initial conditions.

The medium-field case ($B = 1$) was investigated in more detail by running with a first-order integration scheme and using ATHENA++ version 19.0 (Stone et al. 2020) with the same resolution

and static mesh-refinement. Results are plotted in Fig. 5 using PION with the HLLD solver and the second-order scheme (a), the HLLD solver and the first-order scheme (b) and using ATHENA++ (c). Both codes show basically the same result, with small differences because ATHENA++ uses a different integration scheme, especially with regard to integration of the magnetic field. The first-order scheme is more diffusive for all waves, and so features are not as sharp, but the imprint of the refinement boundary remains.

The error appears to be related to an inconsistency introduced in the discretized equations when an oblique shock crosses the refinement boundary, for strongly magnetized plasma with $\beta \lesssim 1$. Significant




Figure 3. MHD Blastwave test calculation in 2D using a uniform grid with 256×384 cells, calculated with PION, for a background magnetic field of strength $B = 0.1$ (a), $B = 1$ (b), and $B = 10$ (c), at an angle of 45° to the positive x -axis. Gas density is plotted at $t = 0.2$ using the indicated linear colour scale. Contours of density are plotted on a linear scale starting from $\rho = 0$ separated by $\Delta\rho = 0.2$.




Figure 4. As Fig. 3, but using a nested grid with two levels, centred on the origin.

effort was made to characterize and eliminate the issue, but no satisfactory solution was found. The features are basically the same whether one uses ideal MHD, ideal MHD with Powell source terms, or ideal MHD with the GLM-MHD divergence cleaning method, although there are small differences in each case. Removing the BC89 flux correction also does not remove the error (or change the solution to any great extent). It is worth noting that ATHENA++ uses a constrained transport scheme to eliminate $\nabla \cdot \mathbf{B}$ (Gardiner & Stone 2005), completely different from the methods used here, and so the issue is not caused by the divergence-cleaning implementation.

Features introduced to the flow by waves crossing refinement boundaries are also discussed in Stone et al. (2020, figs 38 and 39), where a simulation of the relativistic and magnetized Kelvin–Helmholtz instability is run with a uniform grid and with AMR. There are noticeable differences, with the uniform-grid simulation showing much smoother and more symmetric flow. It is unavoidable that refinement boundaries introduce some numerical errors to the solution, but the results in Fig. 5 appear to be a worst-case scenario

in terms of the refinement boundary having an effect on the overall solution. In particular, the results presented below for 3D simulation of magnetized bow shocks and H II region expansion have almost indiscernible artefacts in the flow variables at refinement boundaries, even though in some cases $\beta \sim 1$ in the post-shock medium.

3.3 Expansion of a D-type ionization front

The accuracy of PION in tracking ionization fronts propagating at various speeds from D-type to R-type across a uniform grid was presented in Mackey (2012). The implementation in the nested-grid is very similar, in particular the calculation of optical depths and time-stepping restrictions, applied on a level-by-level basis. In this section, we calculate the D-type expansion of an H II region using the parameters and initial conditions of the ‘Early phase’ calculation of Bisbas et al. (2015).

A source of Lyman-continuum photons emits at a rate $Q_0 = 10^{49} \text{ s}^{-1}$ from the origin, into a uniform neutral ISM of density




Figure 5. The MHD Blastwave test calculation in 2D using static mesh-refinement and two levels, for the case $B = 1$, with the finer level centred on the origin and extending to $x = \pm 0.25$ and $y = \pm 0.375$ calculated using (a) PION with the second-order scheme, (b) PION with the first-order scheme, and (c) ATHENA++ version 19.0. Gas density is plotted at $t = 0.2$ using the same colour scale, as indicated.

$\rho_0 = 5.21 \times 10^{-21} \text{ g cm}^{-3}$ composed purely of hydrogen. The gas has a two-temperature-isothermal equation of state, where neutral gas has temperature $T_0 = 100 \text{ K}$ and sound speed $c_0 \approx 0.91 \text{ km s}^{-1}$, and ionized gas has temperature $T_i = 10^4 \text{ K}$ and sound speed $c_i \approx 12.85 \text{ km s}^{-1}$, with the temperature in partially ionized cells linearly interpolated between these values. The Strömgen radius is $R_s = 0.314 \text{ pc}$ ($0.97 \times 10^{18} \text{ cm}$), and the stagnation radius (where the H II region is in pressure equilibrium with the undisturbed ISM) is $R_{\text{stag}} = (c_i/c_0)^{4/3} R_s = 10.72 \text{ pc}$, approximately 34 times larger.

The early-phase test calculation in Bisbas et al. (2015) was evolved for 0.141 Myr, on a grid that extends to $4R_s \approx 3.9 \times 10^{18} \text{ cm}$ in each dimension. We evolve this solution out to a larger extent of $8 \times 10^{18} \text{ cm}$ in each dimension, so that we can test the adaptive resolution effectively (the D-type expansion only begins at $r \approx R_s$ and the effects of numerical resolution only become clear once a shock and swept-up shell can form). We calculate the radius of the ionization front, R_{IF} , from the ionized volume as follows

$$R_{\text{IF}} = \left(\frac{3}{4\pi} \sum_i y_i(\text{H}^+) V_i \right)^{1/3}, \quad (5)$$

where the sum is over all cells, i , on the domain with cell volume V_i that have H^+ fraction $y_i(\text{H}^+) > 0.01$. This allows a consistent solution even when the shocked shell becomes distorted in multidimensional calculations (cf. Bisbas et al. 2015).

The results for a series of 1D calculations with different spatial resolution are shown in Fig. 6, including the relative difference between the low resolution calculations and the calculation with 8192 cells. Here, we show results using uniform grids with 32, 64, 128, 256, and 8192 cells on the domain $r \in [0, 8 \times 10^{18}] \text{ cm}$. These are compared with the ‘thick-shell’ solution of Williams et al. (2018), which was found to be an excellent analytic solution for the early phase of expansion. Both the ionization front and shock front are effectively discontinuities, for which the order of accuracy is reduced to first order by the slope limiter, and so the relative error of the solution improves approximately proportional to the resolution.

Fig. 7 shows the same information but for 2D simulations with uniform and nested grids with up to three levels of refinement and




Figure 6. Expansion of an H II region simulated in 1D with spherical symmetry, showing how the accuracy improves with increasing numerical resolution. Upper panel: the radius of the ionization front as a function of time, for different resolutions, compared with the ‘thick-shell’ solution of Williams et al. (2018). Lower panel: the relative difference between the radius for a given resolution and the radius when run with 8192 grid cells (absolute value). Line styles are the same as in the upper panel.

different grid resolutions (per level) from 32^2 to 128^2 . The H II region crosses the finest level-boundary at $t \approx 0.05 \text{ Myr}$, and the second level-boundary at $t \approx 0.15 \text{ Myr}$ (shown by the cyan dotted lines). Again the solutions are compared with the 1D calculation using 8192 cells. The accuracy of the 2D uniform-grid solutions is almost the same as in 1D with the same resolution, as expected. For simulations with refined grids, the solution is always better than the uniform-grid solution at the same resolution, but approaches this solution at large radius. When the ionization front is within a refined grid, the accuracy of the solution is comparable to that of a uniform grid with the same cell-diameter. After the ionization front crosses a refinement boundary, the solution accuracy begins to degrade to that of the equivalent uniform grid with the coarser resolution. The 3D results, simulated on one octant with reflection symmetry, are indistinguishable from 2D calculations at the same resolution, and are therefore not shown.

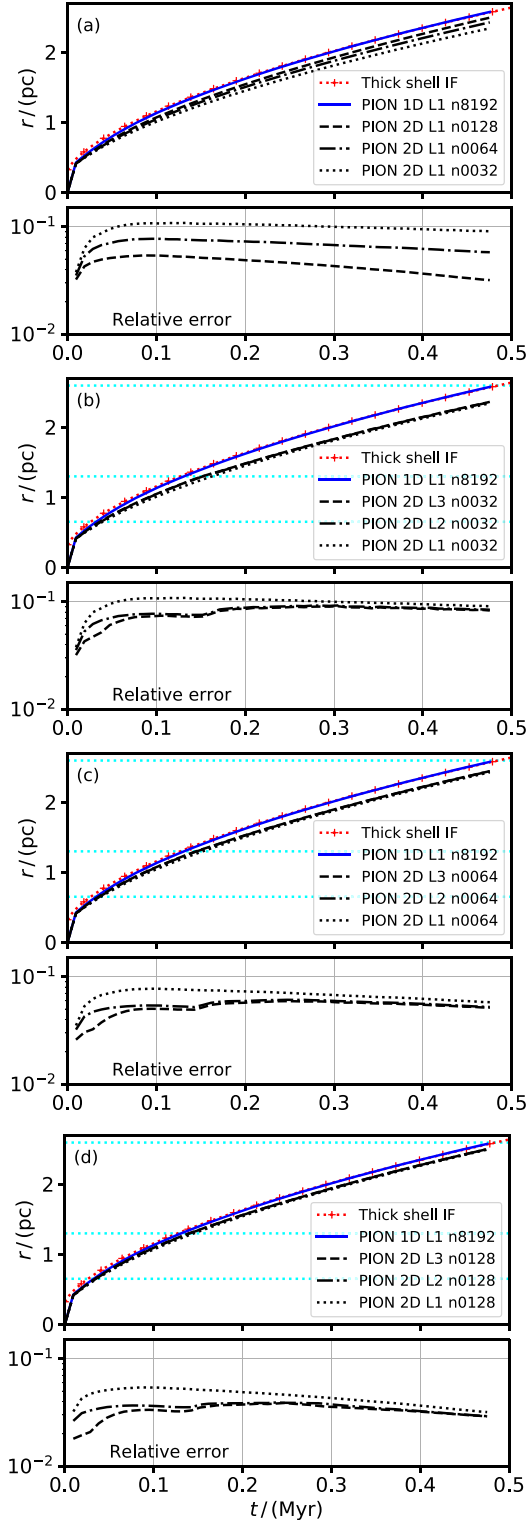


Figure 7. Each panel is as for Fig. 6, except now the calculation is for a 2D simulation of one quadrant with different grid resolutions and numbers of refinement levels. (a) uniform-grid results with different resolutions. (b) grid resolution 32^2 with 1, 2, and 3 refinement levels. Here, the cyan horizontal lines show the boundaries of the three levels. (c) grid resolution 64^2 with 1, 2, and 3 refinement levels. (d) grid resolution 128^2 with 1, 2, and 3 refinement levels.




Figure 8. (a) Maximum shell density as a function of time for 1D simulations of D-type expansion of an H II region, run with differing spatial resolutions from 32 cells to 8192 cells, as indicated in the legend. (b) As above, but with 2D results overlaid in red, for nested grids of 32^2 cells with 1, 2, and 3 levels, as indicated in the legend.

In all cases, the error in ionization-front position is about three cell-diameters, comparable to the numerical resolution of the scheme. This error arises because the shocked shell must be a few cells thick in order to resolve the shock and ionization front, whereas the physical shell thickness is such that the shell is unresolved at all radii for the low-resolution multidimensional simulations shown here. This is shown in Fig. 8(a), where the maximum overdensity in the shocked shell is plotted as a function of time for various different grid resolutions in 1D simulations, from 32 cells to 8192 cells. This is a reasonable proxy for whether or not the shell is numerically resolved, although not sufficient to show that all properties of the shell are correct (e.g. thickness). The simulations using grids with 32 and 64 cells are clearly unresolved at all radii, whereas the grid with 128 cells approaches the correct peak overdensity for $t > 0.4$ Myr, but is increasingly underresolved at earlier times. The overdensity decreases sharply for the high-resolution calculation at the last snapshot because the shocked shell is starting to leave the domain. In Fig. 8(b), 2D results for grids with 32^2 cells and 1, 2, and 3 refinement levels are shown. The results are slightly better than the 1D peak overdensity for an equivalent resolution, but at no stage is the shell resolved.

These results show that the accuracy of the expansion of H II regions for multidimensional simulations with mesh refinement is comparable to that of the equivalent 1D simulation with the same spatial resolution. This means that the dynamics of expanding nebulae can be resolved to approximately the same degree at all stages of expansion. This has advantages for certain applications

Table 1. Stellar wind and ISM parameters for the 3D MHD simulation of a bow shock in section 4.1.

Parameter	Value
Stellar mass-loss rate, \dot{M}	$10^{-7} M_{\odot} \text{ yr}^{-1}$
Stellar wind terminal velocity, v_{∞}	2000 km s^{-1}
Stellar surface rotation (equator), v_{rot}	100 km s^{-1}
Surface split-monopole field strength, $ \mathbf{B} $	10 G
Surface temperature, T_{eff}	35 000 K
ISM density, ρ_0	$2.0 \times 10^{-24} \text{ g cm}^{-3}$
ISM pressure, p_g	$2.9 \times 10^{-12} \text{ dyne cm}^{-2}$
ISM velocity, \mathbf{v}	$[-30, 0, 0] \text{ km s}^{-1}$
ISM B-field, \mathbf{B}_0	$[4, 1, 1] \times 10^{-6} \text{ G}$

such as expanding WR nebulae (Freyer et al. 2006) and Planetary Nebulae.

4 APPLICATIONS TO STELLAR-WIND BUBBLES

Four examples are presented here of simulations of winds from massive stars: a constant wind from an O star driving a bow shock, the nebula produced by an RSG that evolves on a blue loop and spins up to critical rotation, the nebula produced by an RSG evolving to a WR star when it loses its envelope, and the wind–wind collision of two stars in close proximity to each other. The simulations are mainly chosen for ease of comparison with the previous literature on these topics.

4.1 Stellar-wind bubble in 3D with MHD

Here, we introduce the standard wind module in PION, using a spherically symmetric, constant (in time), hypersonic wind from a slowly rotating star (wind type 1 from Section 2.1). We demonstrate for the first time with PION, the implementation of a magnetized wind from a rotating star, a preliminary version of which was presented in Mackey, Green & Moutzouri (2020).

4.1.1 Simulation setup and initial conditions

We set up a 3D MHD simulation of the bow shock produced by a star moving with $v_{*} = 30 \text{ km s}^{-1}$ through the diffuse ISM, with three grid levels. The parameters of the stellar wind and the ISM are in Table 1, and are typical of a late O star of mass $M \approx 30 M_{\odot}$. The ISM pressure corresponds to a gas temperature of $T \approx 7800 \text{ K}$, appropriate for photoionized gas. The standoff distance of the bow shock is defined by

$$R_{\text{SO}} \equiv \sqrt{\frac{\dot{M} v_{\infty}}{4\pi\rho_0(v_{*}^2 + c_s^2)}}, \quad (6)$$

where \dot{M} is the mass-loss rate of the star, v_{∞} is the terminal wind velocity, ρ_0 is the background ISM density, v_{*} is the space velocity of the star, and $c_s \equiv \sqrt{\gamma p_g / \rho_0}$ is the sound speed of the background medium (γ is the adiabatic index of the gas). For these parameters, $R_{\text{SO}} \approx 0.70 \text{ pc}$, and we expect to find the wind termination shock with this separation from the star. The location of the contact discontinuity and the forward shock depend on the compressibility of the gas through the shocks and on radiative cooling efficiency (Scherer et al. 2020).

The radiative heating and cooling prescription follows Green et al. (2019) and the shocked wind is almost adiabatic whereas the shocked

ISM is effectively isothermal. The ISM is assumed to be fully ionized, and to consist of hydrogen with mass fraction 0.714, and helium with abundance one tenth that of H by number, and Solar metallicity (cf. Green et al. 2019). There is no non-equilibrium chemistry, and cooling flag 8 is used, corresponding to photoionized gas that is heated by photoionizations. The heating rate is the product of the local recombination rate $n_e n_H \alpha_B$ and a heating per photoionization of 5 eV, appropriate for an O star. The recombination rate, α_B , is from Hummer (1994). Cooling is the sum of Bremsstrahlung assuming H and He are fully ionized (Rybicki & Lightman 1979; Hummer 1994), and metal-line cooling. For metal-lines, we take the maximum of the Wiersma, Schaye & Smith (2009) collisional ionization-equilibrium (CIE) cooling curve (metals only) and equation (A9) from Henney et al. (2009) for forbidden-line cooling from photoionized CNO ions, which would not be present in CIE because the relevant elements would be at a lower ionization stage at 10^4 K .

The simulation was initialized with a coarse grid of 128^3 grid cells and size $2.4576 \times 10^{19} \text{ cm}$ ($\approx 8 \text{ pc}$) in each dimension (each cell has diameter $\Delta x = 1.92 \times 10^{17} \text{ cm}$). The simulation extents in the x -direction are $x \in [-19.432 \times 10^{18}, 5.144 \times 10^{18}] \text{ cm}$, and $\{y, z\} \in [-12.288 \times 10^{18}, 12.288 \times 10^{18}] \text{ cm}$; the focal point of the nested grids is at $[5.144 \times 10^{18}, 0, 0] \text{ cm}$; and the star is at the origin. We added two levels of refinement to the coarse grid, giving a finest level cell-diameter $\Delta x = 4.8 \times 10^{16} \text{ cm}$. We set the wind inner boundary radius to $9.6 \times 10^{17} \text{ cm}$, corresponding to 20 grid cells on the most refined level.

We used the standard second-order-accurate integration scheme (in time and space) described above, together with the HLL MHD Riemann solver, to evolve the simulation for 10^{13} s , about 1.2 times the crossing time for the star to cross the simulation domain, and nearly 14 times the dynamical time-scale of the bowshock ($\tau_{\text{dyn}} = R_{\text{SO}}/v_{*} \approx 7.2 \times 10^{11} \text{ s}$). The simulation takes about 3 300 CPU-hours to run to completion. For reference, a simulation with 256^3 and three levels would take about $16 \times$ longer (50 000 CPU-h) and 512^3 with three levels would take $\approx 800\,000$ CPU-h.

4.1.2 Results

Results at $t = 10^{13} \text{ s}$ are plotted in Fig. 9, with gas density in panel (a) and the magnitude of the magnetic field in panel (b). A slice through the 3D domain at $y = 0$ is shown, with the star at the origin. The inner 75 per cent of the wind boundary region is set to have very low density, hence the dark circle in panel (a) around the origin. The typical features of the bow shock are seen: the termination shock of the wind, closer to the star in the upstream direction because of the asymmetric ram pressure of the ISM; the strong contact discontinuity, where density changes by a factor of $\gtrsim 10^3$ and the shocked ISM in the bow-shaped arc. The rotation and magnetic axis of the star is \hat{z} ; this has no effect on the gas density profile because the field is relatively weak compared with the ram pressure of the wind, but it can be seen clearly in panel (b).

The magnetic structure of the wind is the classical Parker (1958) Spiral, where the stellar field is wound up by the star’s rotation and drops off as $|\mathbf{B}| \propto r^{-1}$ near the equator and $|\mathbf{B}| \propto r^{-2}$ at the poles. The field switches direction at the equator, producing a current sheet similar to the Heliospheric Current Sheet in the Solar wind. Downstream, the current sheet is swept into the wake behind the star and remains near the $z = 0$ plane, but in the upstream direction, this sheet is swept to large z and back with the flow of the bow shock along the contact discontinuity between wind and ISM. The weakly magnetized polar regions are also swept back in the shocked wind




Figure 9. (a) Gas density, $\log_{10}\{\rho/(\text{g cm}^{-3})\}$, and (b) magnetic field magnitude, $\log_{10}\{|\mathbf{B}|/\text{G}\}$, in the x - z plane through $y = 0$ are plotted on a logarithmic scale as indicated, for a 3D MHD simulation of a bow shock produced by a massive star. The star is at the origin and moving in the $+\hat{x}$ direction. The magnetic axis of the star is \hat{z} , the stellar surface field is $B = 10$ G, and the upstream ISM field is $\mathbf{B}_0 = [4, 1, 1] \times 10^{-6}$ G.

to the wake behind the star. These are the typical features also seen in MHD simulations of the Heliosphere (e.g. Pogorelov, Zank & Ogino 2006), but on larger scales because of the stronger wind. In general, the magnetic and rotation axes may be misaligned (as is the case for the Sun), but this requires a significantly more complicated inner boundary condition (e.g. Pogorelov et al. 2013; Daley-Yates, Stevens & ud-Doula 2019) and is deferred to future work.

The changing resolution is most clearly visible at the contact discontinuity, for which the thickness of the transition layer is mediated by numerical diffusion and therefore scales with the cell diameter. For the HLL solver used here, the transition is spread over approximately 10–15 cells, because the Riemann solver does not contain a contact discontinuity. Shocks are resolved by 2–3 cells, by contrast, and so the effect of resolution is not as obvious. Artefacts such as reflected waves introduced at the resolution boundaries are not visible in Fig. 9 in the way that they were for the MHD blastwave in Fig. 5; the resolution effects are primarily related to the numerical diffusion.

The shocked ISM is asymmetric in the sense that the bowshock is thinner and has higher density in the upper half-plane, although the effect is weak. This is because the ISM magnetic field is closer to the shock normal direction in the upper half-plane than the lower half-plane, and so the magnetic field is less compressed through the forward shock and hence provides less pressure support. The angle between the shock normal and the star’s velocity vector is also larger in the upper half-plane; a geometric measurement of the symmetry axis of the bow shock would therefore be a (in this case, only slightly) biased estimator of the relative velocity between star and ISM.

4.1.3 Stronger stellar magnetic field

A simulation was also run with a stellar magnetic field 10 times stronger, and the results are plotted in Fig. 10 in the same way as for Fig. 9. Such a surface field (100 G) is allowed by observational

upper limits for most O stars (Fossati et al. 2015), although here the Alfvénic Mach number of the wind,

$$\mathcal{M}_A \equiv v_\infty \sqrt{4\pi\rho} / |\mathbf{B}|, \quad (7)$$

has a value $\mathcal{M}_A \approx 10$ (in the freely expanding wind this is independent of radius near the equator because both $\sqrt{\rho}$ and $|\mathbf{B}|$ decrease as r^{-1}). An Alfvénic Mach number much smaller than this would lead to a non-spherically symmetric wind and would require a more complicated inner boundary condition.

Here, there is some accumulation of wind material at the equatorial current sheet (panel a of Fig. 10), on account of the reduced magnetic pressure in this region compared with the neighbouring regions. The contact discontinuity also has a feature near $z = 0$ in the upstream direction, apparently from the sweeping of the current sheet to the upper half-plane. This could be related to similar features seen in the Heliosphere for ideal MHD simulations (Washimi & Tanaka 2001), for which a deeper understanding requires a multifluid description of the flow and/or kinetic theory (Pogorelov et al. 2006). This shows approximately where we expect to reach the limits of applicability of our imposed boundary condition and the single-fluid approximation, i.e. as long as $\mathcal{M}_A \gtrsim 10$ the wind prescription is reasonable.

4.2 Time-varying stellar wind for a rotating star

Chita et al. (2008, hereafter CVL08) studied the formation of ring nebulae around blue supergiants (BSGs) as a result of single star evolution. An RSG embarks on a blue loop, spins up to critical rotation as a result of contraction of the envelope, ejects an equatorially enhanced wind, resulting in a ring nebula around a BSG. We use their work as a test calculation against which to compare results obtained with PION using the latitude-dependent and time-varying wind module (type 2 in Section 2.1). For this test calculation, we use the stellar evolution model F12B calculated by Heger & Langer (2000) and whose circumstellar medium was simulated by CVL08.




Figure 10. As Fig. 9, but for a simulation with a 10 times stronger stellar surface field of $B = 100$ G.




Figure 11. Wind properties for the simulation in Section 4.2: (a) Mass-loss rate, \dot{M} , and Effective Temperature, T_{eff} , and (b) rotation velocity, v_{rot} , critical velocity, v_{crit} , and wind terminal velocity, v_{∞} , as a function of time (since the birth of the star). The RSG phase begins at 18.9 Myr, the phase of critical rotation at 20.035 Myr as the star evolves to hotter surface temperatures on a blue loop.

4.2.1 Stellar evolution calculation and initial conditions

Fig. 11 shows the wind evolution for the evolutionary track F12B, starting at the end of the main sequence at 18.8 Myr, through an RSG phase lasting just over 1 Myr, a blue loop lasting ≈ 0.5 Myr and finishing with a second RSG phase lasting ≈ 0.25 Myr. The beginning of the blue loop is marked by a spike in mass-loss rate driven by the star reaching the Ω -limit as it contracts.

The circumstellar medium was first simulated in 1D with spherical symmetry, up to a point mid-way through the RSG phase, at $t \approx 19.5$ Myr. This is plotted in Fig. 12, where we show H number density, temperature, radial velocity, and H ionization fraction, $y(\text{H}^+)$. This calculation was performed on a simulation domain extending from




Figure 12. Circumstellar medium produced by the F12B stellar evolution calculation in a 1D spherically symmetric R-HD simulation, with the snapshot taken midway through the RSG phase of evolution. This snapshot is mapped on to the 2D grid and used as the initial conditions for modelling the later blue loop.

the origin to $r \approx 52$ pc, resolved by 4096 uniformly spaced grid-cells. The shocked RSG wind forms a diffuse shell at $r \approx 1$ pc with $n_{\text{H}} \sim 1 \text{ cm}^{-3}$, confined externally by a hot and low-density wind bubble from the main-sequence phase. This simulation snapshot was then mapped on to a 2D cylindrical grid with 10 refinement levels and 512×256 cells on each level. The coarsest grid extends to $z \in [-38.9, 38.9]$ pc and $R \in [0, 38.9]$ pc with rotational symmetry in the azimuthal coordinate, and the inner wind boundary has a radius of 0.025 pc (84 grid cells on the finest level). The radial resolution on the nested grid is comparable to that used by CVL08: they used 900 zones to $r = 2$ pc, and we have $256 + 5 \times 128 = 896$ cells to the same radius. The angular resolution of CVL08 is superior at the wind inflow boundary (they use 200 cells for $\theta = \pi/2$ and we have 132, i.e. our angular resolution at the boundary is 0:68).

Photoionization plus gas heating and cooling is solved using the ‘MPv3’ microphysics module that was used for modelling H II regions in Mackey et al. (2013) and photoionization-confined shells around RSGs in Szécsi et al. (2018). Radiative transfer of ionizing photons is included although the stellar temperature only reaches 11 000 K in the BSG phase and so the EUV output of the star is very weak. Collisional heating from shocks is the main heating process active in the simulation (cf. CVL08).

4.2.2 Results

The 2D simulation presented below starts at 19.5 Myr in the middle of the RSG phase and ends at 20.1 Myr, encompassing most of the RSG phase, the period of rapid rotation, and the first 50 000 yr of the BSG phase. The outer boundary conditions are outflow, but the boundary condition for the RSG wind bubble is effectively the confining thermal pressure of the hot and low-density bubble from the main sequence that was modelled in 1D, and which is not shown in the plots below.

Fig. 13 shows snapshots of results (a) at 19.5 Myr in the middle of the RSG phase, (b) just after the phase of critical rotation at 20.04 Myr, and (c) 20 000 yr after onset of critical rotation, at approximately the same time as the final panel of fig. A.4 of CVL08. There are differences in our results for the structure of the RSG wind, compared with CVL08: the termination shock is approximately at the same radius, but there is significantly more post-shock radiative cooling and compression in the CVL08 calculation. This can be traced back to differences in the wind velocity during the RSG phase: we find $v_\infty = [10 - 15] \text{ km s}^{-1}$ whereas CVL08 calculate $v_\infty = [30 - 40] \text{ km s}^{-1}$. This appears to be related to a typo in equation (2) of Eldridge et al. (2006), where the wind multiplier β_w is erroneously multiplying v_{esc}^2 instead of v_{esc} . This would introduce an error of a factor of up to $\sqrt{8}$ for RSGs which can explain the difference.

This modelling difference means that the RSG wind density is different by the same factor at a given distance from the star, and that the interaction of the BSG wind with the swept-up RSG wind proceeds somewhat differently. Nevertheless, the middle panel of Fig. 13 shows that the early part of the wind–wind interaction is very similar in the two studies. A dense equatorial ring is expanding slowly away from the star, and two polar lobes are expanding rapidly, driven by the (now almost spherical) BSG wind with $v_\infty \approx 350 \text{ km s}^{-1}$, sweeping up the slow wind from the RSG phase. Once $T_{\text{eff}} > 10^4 \text{ K}$, the parameter $\beta_w = 1.3$, and so our wind prescription gives v_∞ larger than that of CVL08 by a factor of $\sqrt{1.3}$. Fig. 11 shows that our v_∞ peaks at $v_\infty \approx 370 \text{ km s}^{-1}$ whereas CVL08 have a peak value of $\approx 320 \text{ km s}^{-1}$; the difference is consistent with $\sqrt{1.3}$ times.

The bottom panel of Fig. 13 can be compared with the lower panels of fig. A.4 of CVL08: the BSG wind has swept up a shell to $r \approx 1 \text{ pc}$ in the polar direction, and the shell is thin, radiative, and weakly unstable. In the equatorial plane, a dense ring has expanded to $r \approx 0.25 \text{ pc}$ and this slowly receding ring creates a bow shock in the BSG wind. The termination shock of the BSG wind is well-separated from the swept-up shell in all directions except the equatorial plane. These results are all consistent with CVL08 except that they find the equatorial ring has expanded almost three times as far in this time. This is probably related to the initial injection velocity of the ring, because at all latitudes the wind velocity is multiplied by the parameter β_w .

4.2.3 Including magnetic fields

A stellar magnetic field can be easily included given a prediction or assumption for the time-variation of the surface field of the star. As a




Figure 13. Density and temperature for the 2D simulation described in Section 4.2: the wind bubble during the RSG phase shortly after mapping the 1D simulation on to the 2D nested grid (top panel), just after the phase of critical rotation (middle), and when the swept up wind reaches the termination shock of the RSG wind (bottom). The left half-plane shows $\log_{10} T$ in K and the right half-plane shows $\log_{10} \rho$ in g cm^{-3} . Not all of the simulation domain is shown (the full domain extends to 39 pc).

simple example, Fig. 14 shows the density and temperature field for a similar simulation, but this time assuming a stellar surface magnetic field of

$$B(t) = \left(\frac{10 R_\odot}{R_\star(t)} \right)^2 10 \text{ G}, \quad (8)$$




Figure 14. As Fig. 13, but for an MHD simulation with a simple prescription for the stellar magnetic field. The left-hand half-plane shows $\log_{10}T$ in K and the right-hand half-plane shows $\log_{10}\rho$ in g cm^{-3} , with colour mapping indicated in the colourbars. Not all of the simulation domain is shown (the full domain extends to 2 pc in all directions).

where $R_*(t)$ is the time-dependent stellar radius. Again, we assume that the surface magnetic field is a split monopole swept into a Parker spiral by the stellar rotation, as in Section 4.1. We assume an ISM magnetic field strength $B_z = 10^{-6}$ G. In this case, we could not map the 1D simulation on to a 2D grid because the stellar and ISM magnetic fields break the spherical symmetry, and so we started the simulation at the beginning of the RSG phase, expanding into a uniform ISM with gas density $\rho_0 = 2.338 \times 10^{-24} \text{ g cm}^{-3}$ and pressure $p_0 = 1.318 \times 10^{-12} \text{ dyne cm}^{-2}$. This sets the stagnation radius of the RSG wind bubble, which we chose to be approximately consistent with CVL08.

This simulation used a much smaller domain as a result of the simpler boundary condition: the largest grid extends to $z \in [-2, 2]$ pc and $R \in [0, 2]$ pc with rotational symmetry in the azimuthal coordinate. We use six refinement levels with 640×320 grid cells in each level, and the inner wind boundary has a radius of 0.025 pc (128 grid cells on the finest level, resolving the angular direction by 402 cells over 180°).

The results are very similar to the R-HD calculation, except that the termination shock of the RSG wind is at slightly smaller radius in the MHD simulation and has been overrun at nearly all latitudes by the BSG wind. While the magnetic field here was deliberately chosen to be sufficiently weak that it has little dynamical impact on the nebula, the field configuration and strength is still useful when predicting non-thermal emission such as synchrotron radiation (cf. del Valle & Pohl 2018).

This module can be used, when coupled with results from stellar evolution calculations of single rotating stars, to study the ring nebulae produced by, e.g. spin-up to critical rotation due to stellar contraction, thought to be a potential explanation for the polar caps of Sher 25 and other BSGs (Gvaramadze et al. 2015, CVL08) and for the bipolar structure of nebulae around some LBVs (Langer et al. 1999). The presented R-HD calculation takes ≈ 9000 CPU-h running on 32 cores; an equivalent high-resolution 3D simulation would take $\gtrsim 10^6$ CPU-h. The runtime could be significantly reduced by mapping from 1D at the end of the RSG phase instead of at the middle. The




Figure 15. Evolution of stellar and wind parameters for the $35 M_\odot$ evolutionary model used in section 4.3, from García-Segura et al. (1996b). The mass-loss rate, effective temperature, and wind terminal-velocity are plotted around the evolutionary transition from RSG to WR.

calculations are more demanding than the constant-wind simulations of Section 4.1 because of the radiative transfer and the requirement for high angular resolution at the wind boundary.

4.3 3D R-HD simulation of an expanding WR nebula

García-Segura et al. (1996b) and Freyer et al. (2006) modelled the circumstellar medium of a $35 M_\odot$ non-rotating star that evolved from main sequence to RSG to WR before explosion, using 2D hydrodynamic and R-HD simulations, respectively. Here, we use the same evolutionary calculation, and focus on the RSG→WR transition, for which the wind and surface temperature properties are plotted in Fig. 15. We use the time-varying stellar wind boundary condition (type 1 from Section 2.1) in a 3D R-HD simulation.

4.3.1 Initial conditions and simulation setup

We simulate the main-sequence and RSG phases using R-HD on a 1D grid with 4096 cells covering 100 pc, with a background ISM density of $\rho = 2.338 \times 10^{-23} \text{ g cm}^{-3}$. This is sufficiently large so that the H II region stays in the simulation domain up to the end of the RSG phase, at $t = 4.7542$ Myr. This density field is mapped on to a 3D grid with 256^3 grid cells on each level, and four refinement levels, using zero-gradient boundary conditions. The coarse grid has a range $\{x, y, z\} \in [-30, 30] \times 10^{18} \text{ cm}$ (approximately ± 10 pc) centred on the star. The finest grid has a cell diameter of $\Delta x \approx 2.93 \times 10^{16} \text{ cm}$, and the wind boundary region has a radius of 20 cells ($r \approx 5.86 \times 10^{17} \text{ cm}$, or ≈ 0.2 pc). In comparison with the 2D simulations of García-Segura et al. (1996b), the finest-grid $\Delta x \approx 0.0095$ pc is comparable to their radial resolution $\Delta r = 0.012$ pc for the ‘slow RSG wind’ case, and slightly lower resolution than the $\Delta x = 0.00625$ pc used by Freyer et al. (2006). In terms of the grid used, Freyer et al. (2006) have comparable geometry to ours, modelling one quadrant of the 2D plane with 125^2 grid points per level, whereas we model the full 3D space with 256^3 cells per level, or 128^3 per octant.

4.3.2 Results

The top panel of Fig. 16 plots \log_{10} of gas density and temperature 10 000 yr after the simulation starts. The RSG wind is being swept

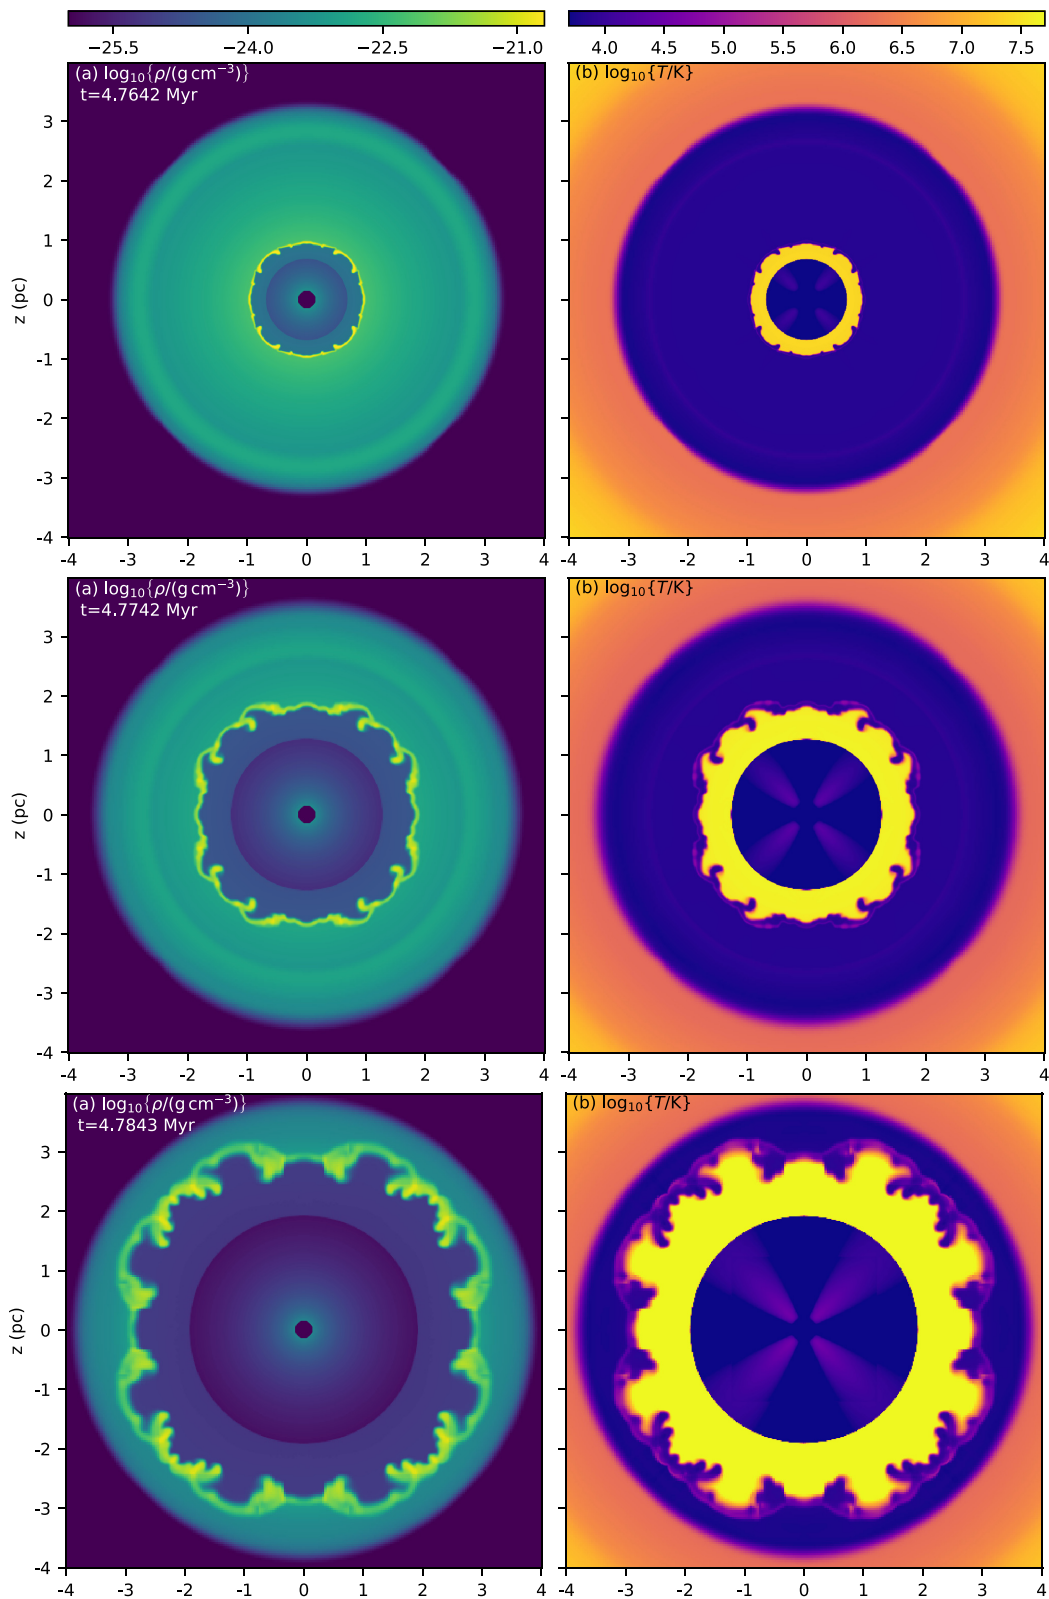


Figure 16. Log of gas density (left-hand panels) and gas temperature (right-hand panels) for a 3D simulation of an expanding WR nebula. In both panels, a slice through the x - z plane at $y = 0$ is shown, with the star at the origin. The axes show the domain in parsecs. The first row shows the CSM 10 000 yr after the RSG \rightarrow WR transition; the second row after 20 000 yr; the third after 30 000 yr.

up into a thin shocked shell that is dynamically unstable, and a strong reverse shock has formed in the WR wind, heating gas to $\gtrsim 10^7$ K. Both the RSG and WR winds are spherically symmetric and so the instabilities are seeded by the grid-scale integration errors. The CSM has already been flash ionized and there is no neutral gas on the domain, except that some of the clumps in the swept-up shell have become optically thick and have partially recombined, but do not cool below $T \sim 10^3$ K.

The hot gas at $r \gtrsim 3$ pc is the relic wind bubble from the main-sequence phase, and this provides the external pressure that previously confined the RSG wind. The RSG wind has been ionized and heated from $T \sim 10^2$ to $\sim 10^4$ K, however, and is now strongly overpressurized, expanding into the surrounding hot medium at approximately its sound speed.

The diagonal features in the temperature plot in the freely expanding wind are integration errors arising from the advection of thermal energy in the strongly kinetic-energy-dominated flow (total energy is conserved by the finite-volume scheme). They do not affect the nebula because the post-shock gas properties are independent of the pre-shock temperature for a large-Mach-number shock.

The second panel shows the gas density and temperature 20 000 yr after the WR wind switches on. The same unstable clumps in the swept-up shell are present, but have grown and have expanded to larger scales. The expansion velocity of the swept-up shell is approximately 1 pc per 10 000 yr (100 km s^{-1}), and so the forward shock is strongly radiative whereas the reverse shock is adiabatic. The third panel plots the results after 30 000 yr, just before the swept-up shell reaches the edge of the RSG wind bubble. The clumpy nature of the shell means that not all directions will break out of the RSG shell at the same time. The symmetry in the solution arises because the hydrodynamic solver is almost perfectly symmetric (to roundoff error), and so the integration errors from the grid discretization are very similar in each octant.

Fig. 17 shows the later evolution as the swept-up shell breaks out of the RSG wind and into the relic main-sequence bubble. This is reminiscent of the simulations presented by Rogers & Pittard (2013), except that their calculations were performed in a dense and turbulent background medium and the WR wind interacted with this interstellar gas and not purely the wind of previous evolutionary phases. Here, the perturbations in the swept-up shell were seeded by the grid geometry rather than a random process, and so the shocked WR wind escapes through eight regularly spaced channels into the low-density surroundings (in this plane; in 3D, there are many channels, but they have a regular spacing), entraining cold RSG wind material as it does so. A more realistic model would introduce clumpy substructure in both the RSG and WR winds, which would break the symmetry and better reflect the reality that winds of massive stars are strongly clumped (Puls, Vink & Najarro 2008), possibly driven by turbulent sub-surface convection (Cantiello et al. 2009; Grassitelli et al. 2015).

4.3.3 Interpretation

Our results are not comparable to Freyer et al. (2006) because they used a wind speed for the RSG phase of $v_\infty = 75 \text{ km s}^{-1}$ whereas our calculation (based on Eldridge et al. 2006, see above) gives a wind speed closer to the ‘slow wind’ calculation with $v_\infty = 15 \text{ km s}^{-1}$ (García-Segura et al. 1996b). We find that the nebula expands to 1 pc in 10 kyr, 2 pc in 20 kyr, 3 pc in 30 kyr, i.e. expanding at $v_{\text{exp}} \approx 100 \text{ km s}^{-1}$. Koo & McKee (1992) calculate the expansion speed of wind-blown bubbles in power-law media, and their equation (3.1), for the case where the reverse shock is adiabatic and the forward

shock is radiative, predicts that the shock radius scales as $R \propto t$ for a constant wind expanding into a density profile $\rho \propto r^{-2}$. The constant prefactor of this equation gives the (constant) expansion velocity as

$$v_{\text{exp}} = \left(\frac{\dot{M}_{\text{wr}} v_{\infty, \text{wr}}^2 v_{\infty, \text{rsg}}}{3 \dot{M}_{\text{rsg}}} \right)^{1/3}. \quad (9)$$

For the WR wind parameters shortly after the transition ($\dot{M}_{\text{wr}} \approx 2.75 \times 10^{-5} M_\odot \text{ yr}^{-1}$ and $v_{\infty, \text{wr}} \approx 1200 \text{ km s}^{-1}$) and RSG wind parameters mid-way through the RSG phase ($\dot{M}_{\text{rsg}} \approx 8 \times 10^{-5} M_\odot \text{ yr}^{-1}$ and $v_{\infty, \text{rsg}} \approx 10 \text{ km s}^{-1}$), this corresponds to $v_{\text{exp}} \approx 120 \text{ km s}^{-1}$, in good agreement with our numerical results. García-Segura et al. (1996b) found somewhat faster expansion (their fig. 7), but the differences are probably not significant. The qualitative appearance of the results are very similar, given that the grid geometry is different and the development of instabilities is not expected to be identical. The formation of clumps in the thin shell, persistent through its expansion, followed by blowout once the edge of the RSG wind bubble is reached, is the same.

This calculation took about 15 000 core-h, run on 32 cores for 20 d. Higher resolution is desirable for better resolving the instability of the swept-up shell, and is required for modelling rotating stars with clumpy winds to resolve the spatial and temporal variations in the wind. Running with 384^3 cells on each level would take ≈ 75 000 core-h, whereas 512^3 would require ≈ 250 000 core-h.

4.4 3D MHD simulation of wind–wind collision

Colliding-wind binary systems are fascinating environments to study astrophysical fluid dynamics, shocks, and instabilities (Stevens et al. 1992; Lamberts, Fromang & Dubus 2011; Parkin et al. 2011; Madura et al. 2013), magnetism (Walder, Folini & Meynet 2012; Kissmann et al. 2016), and particle acceleration (Pittard & Dougherty 2006; Grimaldo et al. 2019; White et al. 2020). In principle, the nested-grid setup of PION is suitable for modelling such systems, although some code improvements would be required, including implementation of orbital motion for stellar-wind sources, radiative heating/cooling routines that are appropriate for the high densities and chemical abundances encountered, and possibly a wind acceleration region for close binary systems. Here, we present a preliminary 3D MHD simulation of wind–wind collision based on the 2D hydrodynamic calculation of the V444 Cyg system by Stevens et al. (1992) using the constant wind module (type 0 in Section 2.1).

V444 Cyg is a well-studied eclipsing binary system consisting of two massive stars with powerful winds, a WR primary (spectral type WN5) and an O-type secondary (spectral type O6). For the modelling of the colliding winds, we took the orbital and stellar wind parameters from Stevens et al. (1992), presented in Table 2. We make a number of simplifications to the system, partly so the setup is comparable with Stevens et al. (1992)

- (i) We inject the winds already at the terminal velocity at the wind-boundary radius, and radiation and gravitational forces are neglected. We set the wind boundary radius for both stars at $r = 2.6 \times 10^{11} \text{ cm}$.
- (ii) The orbital period of the system is 4.2 d, nevertheless during this 3D simulation, the orbital motion is neglected, because of the extra complications this would introduce (the system geometry will change, and must be simulated for much longer to reach a stationary state). Thus, we consider that the stars were at rest in an inertial frame with cylindrical symmetry along the x -axis, except for the stellar magnetic fields which break the symmetry.
- (iii) We use an optically thin cooling function appropriate for photoionized bow shocks (Green et al. 2019), although this is likely




Figure 17. As Fig. 16, but showing results at 35 000 and 45 000 yr after the RSG→WR transition, and zoomed out to show the expanding nebula.

Table 2. The properties of the stars of V444 Cyg system, where primary is a WR star and secondary has spectral type O6, taken from Stevens et al. (1992). The rotation velocities and surface split-monopole magnetic-field strengths are notional, for demonstration of the methods only, and both have axis of symmetry \hat{z} .

Parameter	Primary	Secondary
Mass-loss rate, \dot{M} ($M_{\odot} \text{ yr}^{-1}$)	1.4×10^{-5}	10^{-6}
Terminal wind speed, v_{∞} (km s^{-1})	2000	2000
Surface rotation speed, v_{rot} (km s^{-1})	200	200
Surface split-monopole magnetic field strength, $ \mathbf{B} $ (G)	100	1
Radius of wind boundary (cm)	2.6×10^{11}	2.6×10^{11}
Position of star (10^{12} cm)	$[-1.84, 0, 0]$	$[0.96, 0, 0]$

a crude approximation given the large densities and the hydrogen depletion of the WR wind.

The simulation is run with three grid levels, each with 384^3 cells and with the two nested grids centred on the origin at the centre of the domain. The coarsest grid has $\{x, y, z\} \in [-1.024, 1.024] \times$

10^{13} cm, the next level has $\{x, y, z\} \in [-5.12, 5.12] \times 10^{12}$ cm, and the finest level has $\{x, y, z\} \in [-2.56, 2.56] \times 10^{12}$ cm, with a cell diameter of $\Delta x = 1.333 \times 10^{10}$ cm. Outflow boundary conditions are employed at all boundaries.

We added stellar rotation and split-monopole magnetic fields for each star, such that the magnetic field in the wind will be swept into a Parker spiral at large radius. The field strengths were chosen such that the magnetic field is too weak to affect the dynamics of the unshocked wind, i.e. the Alfvén Mach number of the wind is large in both cases. The stars were set rotating well below critical rotation so that latitude-dependent effects are not expected to be strong.

At the first stage, we run the simulation with adiabatic hydrodynamics, without taking radiative cooling into account. The simulation is run for 70 000 s, which is enough time for a stationary shock structure to form around the stagnation point of the flow. The dynamical time-scale of the wind–wind collision is the stellar separation divided by the wind speed, 14 000 s. The results of the simulation are plotted in Fig. 18 for gas density and magnetic field strength. A stable wind–wind collision has been set up and the weaker wind of the O star gets swept back by the stronger WR wind. The

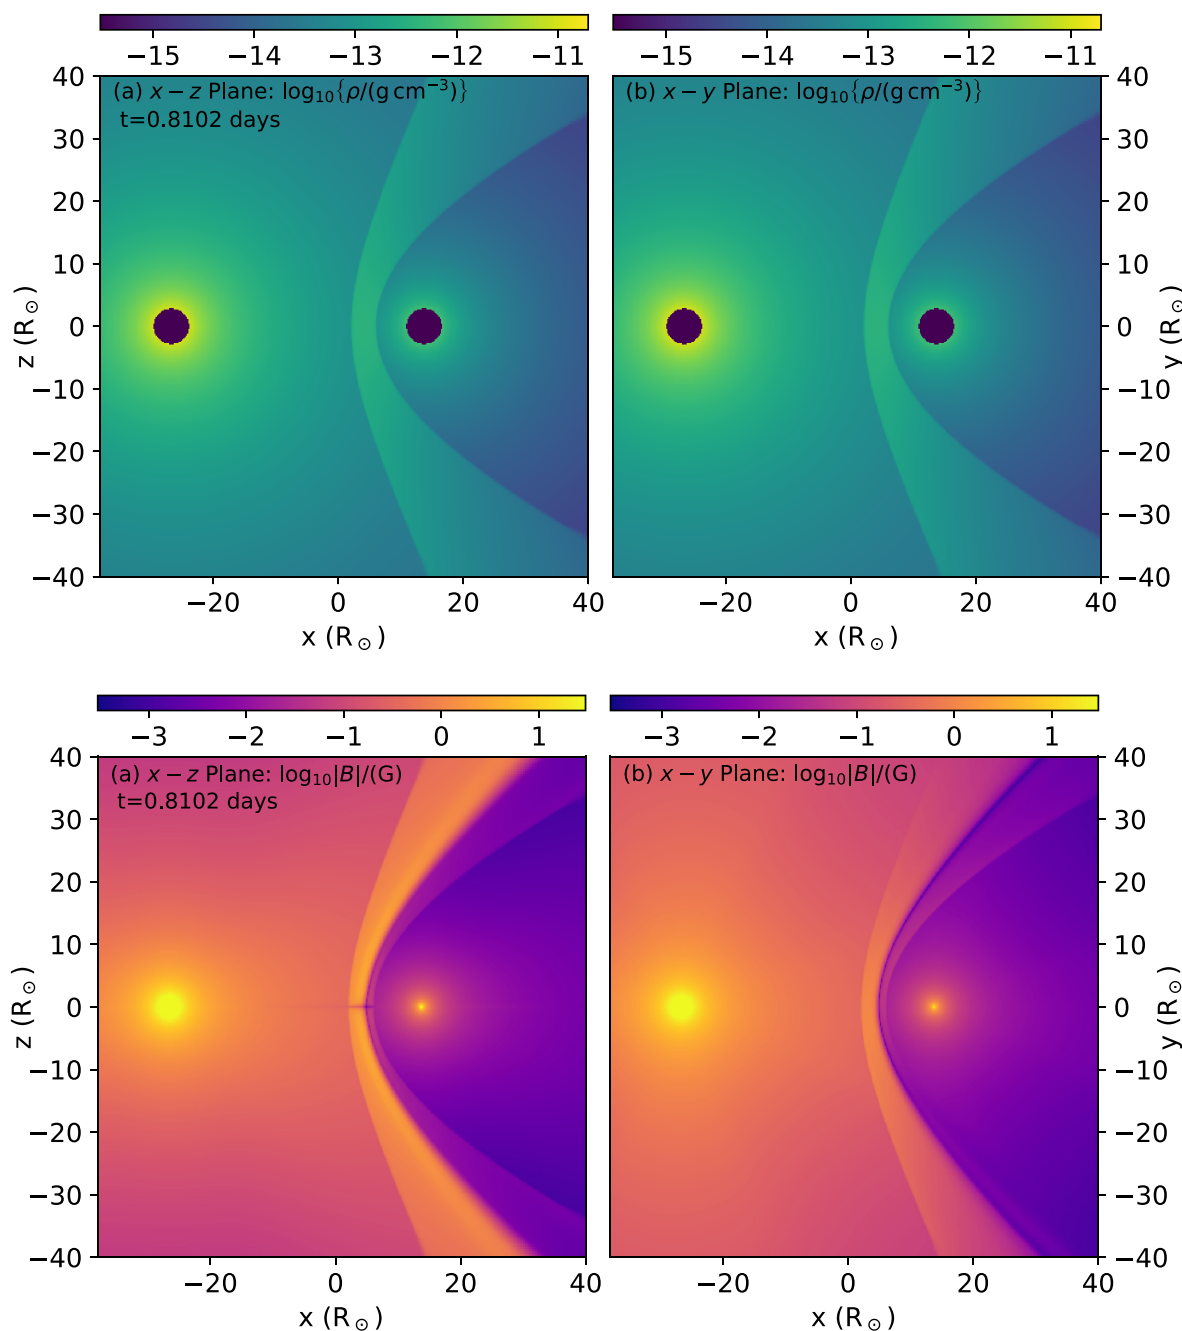


Figure 18. 3D MHD simulation of wind-wind collision from Section 4.4, for the adiabatic case at $t = 70\,000$ s. Above: \log_{10} of gas density in the x - z plane (left-hand panel) and the x - y plane (right-hand panel). Below: \log_{10} of magnetic field strength in the x - z plane (left-hand panel) and the x - y plane (right-hand panel).

collision region, shaped like a bow shock, consist of two shocks and a hot plasma between these shocks.

For the second stage, we continue the simulation including radiative cooling. As noted by (Stevens et al. 1992), for this system the cooling time is comparable to the advection time for the shocked wind, and so the gas cools and can be compressed to very high densities. The wind-collision region is unstable, i.e. the shocked region gets narrower and eventually an oscillatory thin-shell instability arises.

Fig. 19 shows the same plots as Fig. 18, but at a later time after radiative cooling has been switched on for some time. The shocks

oscillate between strongly radiative and weakly radiative near the stagnation point of the flow. This appears superficially similar to the case of overstable radiative shocks with velocity ≈ 150 km s $^{-1}$ (Innes, Giddings & Falle 1987), but is actually arising because the advection time and cooling time for the shocked gas are very similar for this setup (Stevens et al. 1992). Small perturbations in the hydrodynamics can make the difference as to whether a parcel of gas can cool strongly or not before it is advected away from the stagnation point. The features seen in all panels near the stagnation point are transitory and unsteady, with knots and rope-like overdensities forming and advecting away to the domain boundaries. The shocks

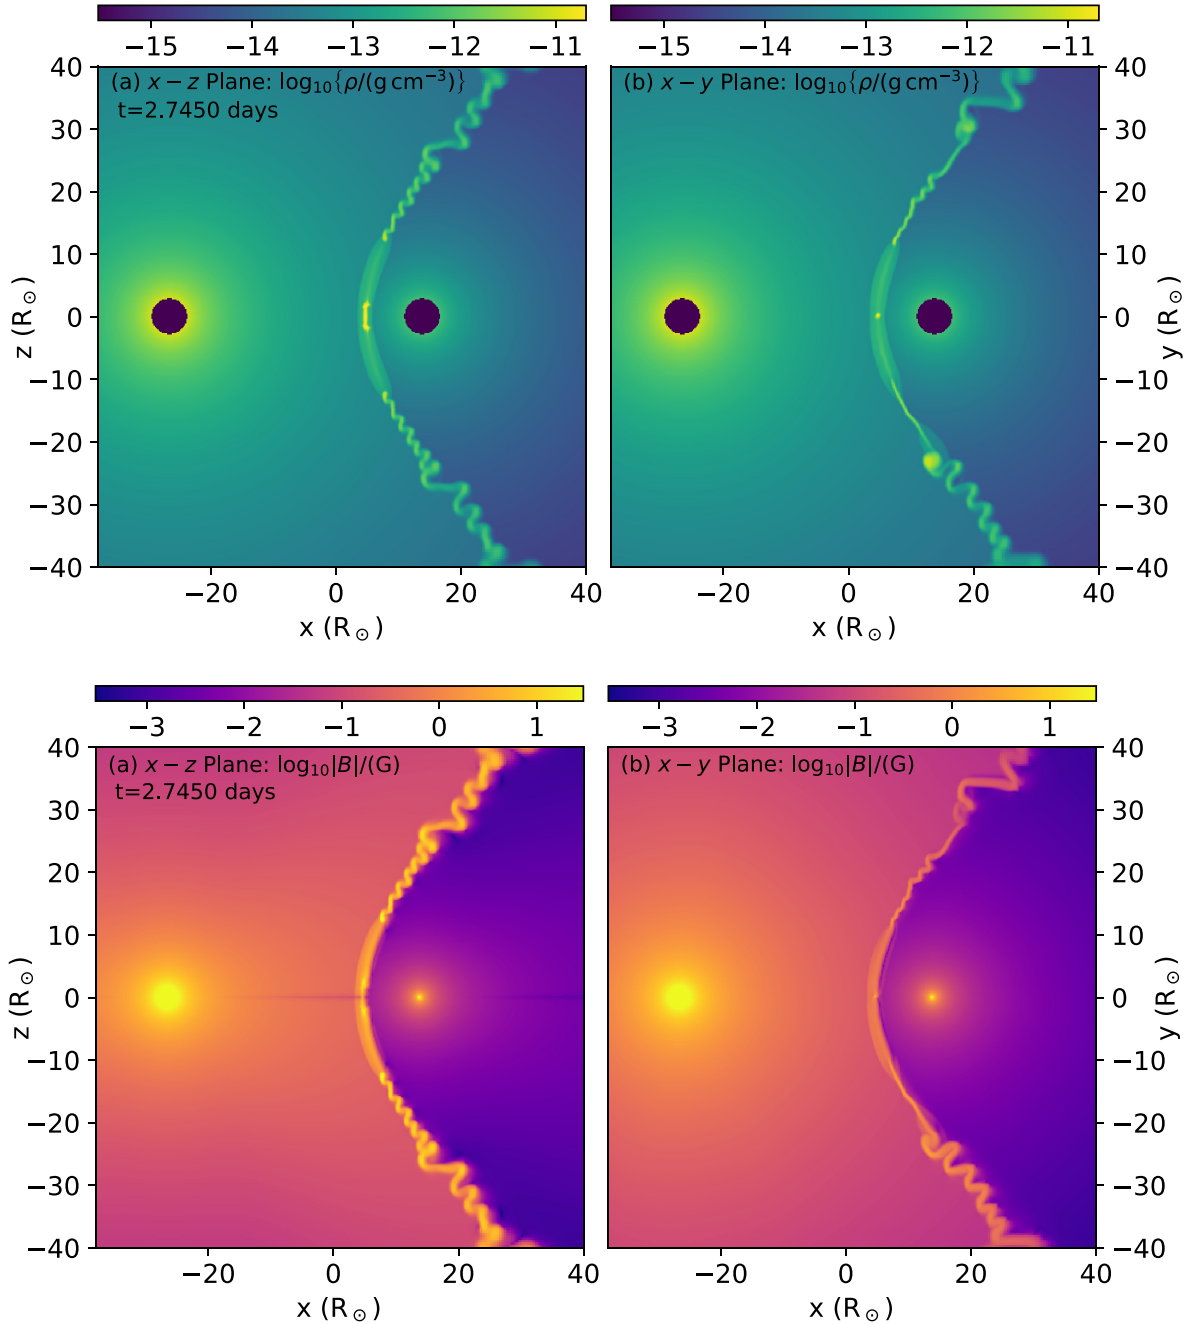


Figure 19. 3D MHD simulation of wind–wind collision from Section 4.4, at $t \approx 2.4 \times 10^5$ or $\approx 1.7 \times 10^5$ s after radiative cooling has been switched on. Above: \log_{10} of gas density in the x – z plane (left-hand panel) and the x – y plane (right-hand panel). Below: \log_{10} of magnetic field strength in the x – z plane (left-hand panel) and the x – y plane (right-hand panel).

far from the symmetry axis of the flow are strongly radiative and do not show this oscillatory behaviour. This is because the shocks are oblique, with smaller Mach number than along the symmetry axis, and so the post-shock temperature is lower and the cooling time is shorter (cooling time has a maximum at $T \approx 2 \times 10^7$ K and decreases for both lower and higher temperatures).

Gas compression factors of between 10 and 100 are achieved in the radiative shocks (limited by the grid resolution), with somewhat weaker increase in the magnetic-field strength because only the component perpendicular to the shock normal is compressed. For

the small separation of the two stars, the magnetic fields of the two winds near the stagnation point are still more radial than toroidal, and only the toroidal component is amplified. The magnetic-field amplitude in the x – y plane is less than in the x – z plane because the former contains the equatorial current sheet of both stars.

A detailed investigation of the V444 Cyg system would require the inclusion of orbital motion (cf. Lamberts et al. 2017), the finite size of the stars, their wind-acceleration regions, and perhaps also radiative inhibition (e.g. Parkin et al. 2011). In future work, we will improve the radiative cooling function and implement orbital motion, and use

higher resolution simulations to investigate the MHD properties of shocks in colliding-wind binary systems. The simulation presented here took 10 000 core-h, running with 128 MPI processes.

5 POST-PROCESSING SIMULATION SNAPSHOTS

5.1 PYTHON library for reading and plotting snapshots

A PYTHON library (PYPION¹) has been developed to read PION snapshots into NUMPY arrays for plotting and further analysis. This enables simple post-processing and visualization on all simulations using PYTHON.

PYPION contains two core scripts that contain modular routines that can be used depending on what type (1D/2D/3D) of PION simulation is run. The library also works on simulations with refined grids without modification. An additional script (`PlottingClasses.py`) provides some examples of plots that can be generated from PION data. Most of the figures in this paper showing multidimensional simulations, including e.g. Figs 9, 10, and 16–19, have been produced using this library.

(i) `SiloHeader_data.py` contains a class with methods to open a PION snapshot in SILO format² and read metadata from the header directory. This data includes the axes dimensions, level dimensions, number of levels, simulation time, number of MPI processes.

(ii) `ReadData.py` contains a class with methods for reading the data for a requested variable (e.g. Density) from a SILO file, returning it as a single NUMPY array per level. When PION is run with multiple MPI processes, each process calculates a subdomain of the grid for each snapshot and saves its data under a SILO directory. This class reads each subdomain in turn and adds its data to the correct region in the NUMPY array to form an image.

(iii) `PlottingClasses.py` contains classes to take data from `ReadData.py` and make some commonly used plots using MATPLOTLIB. Several functions have been set up to accommodate different user needs.

5.2 Radiative transfer with TORUS

We have implemented a method for post-processing 3D nested-grid PION simulations with the TORUS Monte Carlo radiative transfer code (Harries et al. 2019). This builds on the previous implementation that post-processed 2D uniform-grid PION simulations with TORUS to make synthetic dust continuum images (Mackey et al. 2016; Gvaramadze et al. 2017; Green et al. 2019). An improvement compared with the method of Green et al. (2019) is that we now use a passive tracer to distinguish wind from ISM, and this is used to set the dust-to-gas ratio to zero in the wind and to 0.01 in the ISM, with a smoothly varying interface region where wind and ISM are mixed. Previously, we used a simple temperature cut to distinguish wind and ISM, which was effective but not as self-consistent as the new treatment.

PION simulation snapshots are converted to FITS format and read into TORUS using a C++ programme, `SILO2FITS`, provided with the PION source code. The TORUS reader maps variables including the density, temperature, and dust-to-gas mass ratio on to the TORUS grid using a bilinear (for 2D models) or trilinear (3D) interpolation. In static mesh-refinement applications, each level of the PION grid

is stored in a separate FITS file and these are read into TORUS sequentially. Any given cell on the TORUS grid is populated by the highest resolution PION data available. The resolution of the TORUS grid is flexible and can reflect the structure of the grid being read in, or adaptively refine according to its own AMR grid criteria (e.g. mass per cell and/or gradients in any quantity), or revert to a uniform mesh.

With the physical parameters of the PION grid read in, appropriate stellar parameters (temperature, radius, location) are added manually to the TORUS input file. With the grid and photon sources set up, the full scope of TORUS functionality is then available (Harries et al. 2019). For example, this means that the thermally decoupled dust-temperature can be computed in a Monte Carlo radiative equilibrium calculation for arbitrary grain composition and size distribution. For this test calculation, we use silicate grains (Draine 2003) with a Mathis, Rumpl & Nordsieck (1977) size distribution from 0.01 to 10 μm .

Fig. 20 shows the PION gas density and temperature in panels (a) and (b) compared with the dust-to-gas ratio (panel c) and the dust temperature, T_D , (panel d) obtained from the TORUS radiative equilibrium calculation. The dust density corresponds very closely to the gas density, except that inside the contact discontinuity the dust density decreases to zero in the inner part of the wind bubble, and so it is not shown. The dust temperature is completely decoupled from the gas temperature because the collisional heating and cooling rates are negligible compared with radiative rates for these diffuse ISM conditions (e.g. Meyer et al. 2014). In the inner part of the bow shock $T_D \approx 30$ K, whereas further out T_D decreases to ≈ 20 K.

The resulting dust emission maps at wavelength 50 μm are shown in Fig. 21 from a range of viewing angles from edge-on to face-on. The classic parsec-scale bow-shock morphology (e.g. Peri et al. 2012; Peri, Benaglia & Isequilla 2015) is seen in the edge-on panels, where limb-brightening makes the shocked shell much brighter than is seen in the face-on panels. The relatively low spatial resolution, combined with low-density ISM and small space velocity of the star, ensures that the bow shock is smooth with no apparent instability at either the contact discontinuity or forward shock. There is a small asymmetry between the upper and lower half-plane, arising because the ISM magnetic field is not parallel to the star's motion, and so the shock compression factor and Mach number is not symmetric about $z = 0$.

6 PERFORMANCE AND PARALLEL SCALING

A number of scaling tests have been performed to assess the performance of PION on HPC systems. All calculations were run at the Irish Centre for High-End Computing (ICHEC) on the supercomputer *Kay*,³ using the cluster nodes each consisting of 2×20 -core 2.4 GHz Intel Xeon Gold 6148 (Skylake) processors with 192 GiB of RAM and a 100 Gbit OmniPath network adaptor.

6.1 Strong scaling for 3D MHD

To test the strong scaling of PION, a 3D simulation of a bow shock with ideal MHD was run, with a resolution of 256^3 grid cells on each level, and with three levels of refinement. Note this calculation has no radiative transfer, so there are no long-range interactions and we expect the scaling to be good up to the point where the number of grid cells being communicated in boundary data is comparable with the the number of grid cells being calculated per MPI process.

¹<https://git.dias.ie/massive-stars-software/pypion>

²<https://wci.llnl.gov/simulation/computer-codes/silo>

³<https://www.ichec.ie/about/infrastructure/kay>

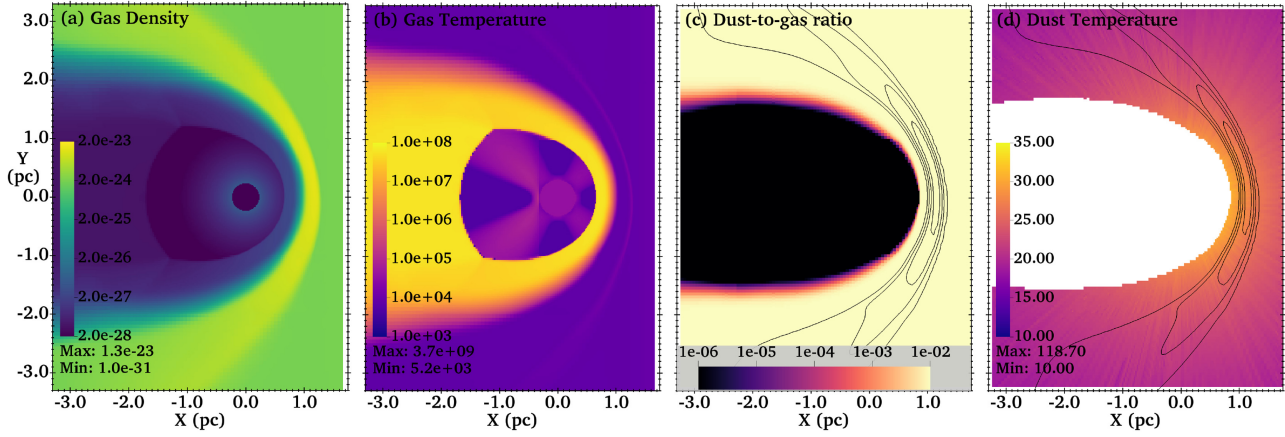


Figure 20. 2D slice through the 3D MHD bow-shock simulation from Section 4.1 for the star with a 10 G surface magnetic field showing (a) gas density in g cm^{-3} , (b) gas temperature in K, (c) dust-to-gas (mass) ratio and (d) dust temperature in K, all on logarithmic colour scales except for dust temperature which is on a linear scale. Dust temperature is calculated using TORUS, and regions with dust-to-gas ratio $< 10^{-4}$ are masked for clarity. In panels (c) and (d) contours of gas density are plotted, linearly spaced from 0 to $10^{-23} \text{ g cm}^{-3}$ in steps of $0.25 \times 10^{-23} \text{ g cm}^{-3}$.




Figure 21. Dust-emission map of the 3D MHD bow-shock simulation from Section 4.1 for the star with a 10 G surface magnetic field, calculated using TORUS. From left- to right-hand and top to bottom, the panels show projections with image normal vector at an angle of 90° , 75° , 60° , 45° , 30° , and 15° with respect to the positive x -axis. The $50 \mu\text{m}$ intensity is plotted in MJy ster^{-1} on a linear scale. The empty region upstream from the bow shock in the upper panels is outside the PION simulation domain.

The star has a mass-loss rate of $\dot{M} = 1.74 \times 10^{-6} M_\odot \text{ yr}^{-1}$, wind velocity $v_\infty = 2500 \text{ km s}^{-1}$, and is placed at the origin of the simulation domain. The star is moving through the ISM at $v_* = 30 \text{ km s}^{-1}$ in the \hat{x} direction, modelled as a flow past the star (which is static on the computational domain), and the uniform ISM has number density is $n_0 = 100 \text{ cm}^{-3}$. The stellar magnetic field is taken to be a split monopole (radial field lines) with a surface field strength of $B = 10 \mu\text{G}$. The interstellar magnetic field is oriented perpendicular

to the star's space velocity, and has a strength $B_z = 25 \mu\text{G}$. The setup is otherwise similar to the simulation in Section 4.1, just with different star/wind and ISM parameters.

We assume that H and He are both singly ionized by the star's radiation field in the full domain, so radiative photoheating balances radiative cooling at a temperature $T \approx 7500 \text{ K}$, following the same heating/cooling routines as Green et al. (2019). The simulation was evolved to $t = 2.37 \times 10^{12} \text{ s}$, about 25 per cent of the time required

Table 3. Strong scaling of PION for simulation of a bow shock in 3D with three levels of refinement and 256^3 grid cells per level, for 2048 time-steps.

N_{proc}	Walltime (s)	Core-h	Speedup	Efficiency
32	18 176	161.6	1.0	1.00
64	10 101	179.6	1.8	0.90
128	5 791	205.9	3.14	0.78
256	3 272	232.6	5.56	0.69
512	1 989	282.9	9.14	0.57
1024	1 868	531.3	9.73	0.30




Figure 22. Speedup of PION for a strong scaling test, running a 3D bow-shock simulation with three levels of refinement and 256^3 grid cells per level, for 2048 time-steps. The solid black line shows the attained speedup, the dashed blue line is the ideal case, and the dotted blue line is the best possible speedup taking account of the extra ghost-cell calculations introduced by domain decomposition, but assuming communication overhead is zero. The data are from Table 3.

to reach a stationary state, and a snapshot was saved. This was taken as a starting point for the scaling test, which consisted of continuing the simulation for 2048 time-steps with MPI process counts between $N_{\text{proc}} = 32$ (run on a single node) to $N_{\text{proc}} = 1024$ (on 26 nodes). The results are shown in Table 3 and plotted in Fig. 22, where the speedup, S , is defined as the wall-time to run the calculation on 32 MPI processes, T_{32} , divided by the wall-time to run on N MPI processes, T_N . The ‘ideal’ case is perfect scaling, where doubling the number of MPI processes will decrease the run time by a factor of 2. The ‘best’ case takes account of the extra ghost-cells that must be calculated when the full domain is subdivided into more sub-domains, given that the boundary region is four cells thick, but assumes zero communication overhead. For a 256^3 grid, this is $S = 11.02$ for $N = 512$, and the ideal value is $S = 16$. Compared with a simulation with 32 MPI processes, a calculation with 512 MPI processes is still 57 per cent efficient, and the code speeds up by a factor of 9.14 compared with a theoretical best attainable value of 11.02.

Although this strong scaling is good, and allows us to run 3D MHD simulations efficiently on hundreds of cores, we have not tested the scaling to $\gtrsim 10^4$ cores. The ratio of ghost cells to grid cells increases strongly as the number of MPI processes increases and the subdomains assigned to each MPI process get smaller. From Fig. 22, it is clear that there could be significant gain by switching to a hybrid parallelization scheme, using multithreading to reduce the number




Figure 23. Strong scaling of PION for 3D R-HD with four levels of refinement and 256^3 grid cells per level, run for 1536 time-steps. Curves are as for Fig. 22, except that the blue dotted line now shows a scaling of $N_{\text{proc}}^{2/3}$, i.e. the expected performance of the 3D radiative transfer algorithm from uniform-grid calculations in Mackey (2012).

of MPI processes per node. We plan to implement this for the next release of PION.

6.2 Strong scaling for 3D R-HD

Radiative transfer scales less well than hydrodynamics because we work in the limit where the speed of light is infinite, and so there are long-range interactions between each cell and each radiation source that must be calculated by tracing rays from one sub-domain to the next. The uniform-grid version of PION had strong-scaling speedup of $S \propto N_{\text{proc}}^{0.5}$ for 2D calculations and $S \propto N_{\text{proc}}^{2/3}$ for 3D calculations for the radiative transfer part of the calculation (Mackey 2012). For the 3D nested grid, we consider an expanding WR nebulae from Section 4.3, taking an initial snapshot from $t = 4.7638$ Myr, about 10 000 yr after the stellar transition from RSG to WR. The simulation domain has four levels of refinement centred on the stellar source at the origin, with 256^3 grid cells on each level.

The simulation was run for 1536 time-steps (on the finest level), corresponding to about 2400 yr of evolution, and the speedup is plotted in Fig. 23. The scaling is significantly worse for R-HD than for the MHD simulations without radiative transfer and non-equilibrium ionization. Increasing the core-count by a factor of 8 already reduces the efficiency by more than 50 per cent, and running this simulation on 1024 cores requires four times more core-h than running on 32 cores. The slope of the scaling plot is approximately constant except for the jump from 32 to 64 cores, corresponding to switching from calculating on a single node to multiple nodes with slower communication. Similar difficulties obtaining good scaling for R-HD have been reported previously (e.g. Wise & Abel 2011), although some innovative algorithms have improved scaling to larger numbers of cores (Rosen et al. 2017).

6.3 Weak scaling for 3D R-HD

Here, we consider the expansion of an H II region and wind bubble from a massive star into a uniform and static ISM. The medium is dense and the wind is strong, and so the ionization front is trapped by the forward shock driven by the expanding wind bubble. We take

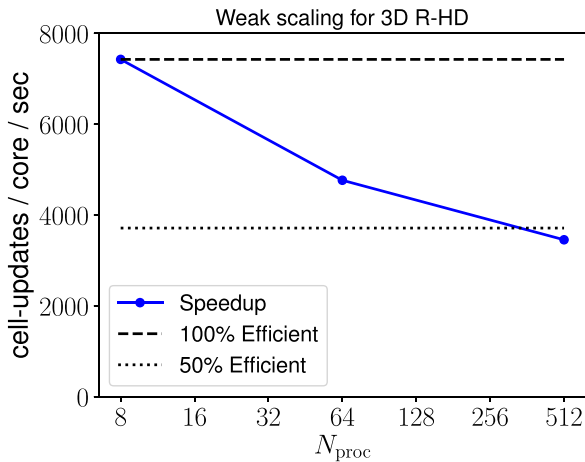


Figure 24. Weak scaling of PION for 3D radiation-hydrodynamics with four levels of refinement and 64^3 grid cells per MPI process, run for ≈ 4.5 h walltime.

the result from a 1D calculation with PION and map it on to a 3D grid with four refinement levels each with 128^3 , 256^3 , or 512^3 cells. The weak scaling is tested by running calculations where the number of grid cells per core is constant, so the 128^3 simulation is run on 8 cores, the 256^3 on 64 cores, and the 512^3 on 512 cores. Each core therefore computes a subdomain of 64^3 cells in all three cases.

The results are plotted in Fig. 24, where we show the number of cell-updates per core per second for simulations with 8, 64, and 512 MPI processes. The efficiency of the code decreases by about 50 per cent when the number of MPI processes increases by a factor of 64. The overall performance of PION for this calculation is not optimal, running at nearly 20 times slower than simple MHD without any radiative transfer or non-equilibrium-ionization calculation. There is certainly scope for improving both the overall performance and the parallel scaling of this algorithm, and the size of simulation that can be run is currently limited by the parallel scaling.

7 CONCLUSIONS

We have presented upgrades to the simulation framework PION for astrophysical fluid dynamics, including the first public release of the source code and associated scripts and post-processing routines. The major upgrades are the implementation of static mesh-refinement, the improved robustness of the MHD solver (including improved divergence cleaning), the implementation of the consistent multispecies advection (sCMA method) for advecting elements across the domain, and the addition of latitude-dependent and magnetized winds from rotating stars following Langer et al. (1999) and Pogorelov et al. (2004).

Test calculations showing advection and the expansion of blast waves and ionization fronts across refinement boundaries have been presented. The advection of a magnetic field loop shows no noticeable artefacts associated with the refinement boundaries. Blast-wave expansion also works very well for hydrodynamics and for weak magnetic fields, but an artefact appears once the magnetic field becomes dynamically important. An imprint of the refinement boundary is apparent in the expanding shocked medium, also when run at first-order accuracy and with different methods for divergence cleaning. Comparing with ATHENA++ (Stone et al. 2020) when run on the same problem, the results are very similar, even though ATHENA++ uses a different integration scheme and a different

algorithm for dealing with magnetic-field divergence errors. Looking at the D-type expansion of an ionization front, PION produces results with static mesh-refinement that are at least as accurate as uniform-grid simulations with equivalent resolution, and with a fraction of the computational cost.

Results were presented for a 3D MHD bow shock produced by the wind of a rotating and magnetized O star moving with 30 km s^{-1} through a uniform ISM, a preliminary version of which was presented in Mackey et al. (2020). The classic features of a Parker (1958) wind were demonstrated: the Parker spiral, equatorial enhancement of the toroidal magnetic field, and the equatorial current sheet. It was shown that for a reasonably strong surface split-monopole magnetic field of 100 G, the magnetic field strength in the shocked wind bubble can be comparable to that in the shocked ISM. For bow shocks, where synchrotron radiation can be detected (e.g. BD + 43 3654 Benaglia et al. 2010), it may be possible to constrain this magnetic-field strength observationally, giving a direct constraint on the stellar magnetic field.

We revisited the calculation of CVL08 of the ring nebula produced when a rotating RSG evolves on a blue loop and reaches critical rotation. Using the latitude-dependent wind prescription of Langer et al. (1999), we showed that PION produces results with a 2D nested grid in cylindrical coordinates that are comparable to the 2D spherical-grid computations of CVL08. We largely confirm their results, although there are some small differences in the details of the hydrodynamic flow. We also demonstrated how an MHD simulation of such a nebula can be calculated by making a simple assumption about the surface magnetic-field strength. A more realistic calculation would estimate the stellar magnetic field strength from the properties of the stellar envelope. The code is efficient enough that 3D simulations of ring nebulae are feasible with reasonable computational resources.

Comparing PION with another classic calculation of a circumstellar nebula, we used a 3D nested-grid to simulate the expansion of a spherically symmetric fast wind from a WR star into the slow wind from its previous RSG phase of evolution. We used the same stellar evolution calculation as the 2D simulations by García-Segura et al. (1996b) and Freyer et al. (2006), and our 3D simulation has comparable spatial resolution to the previous 2D work. Again, our results are comparable to previous work, although there are some differences in the details. We showed that the WR wind bubble expands at almost constant speed as predicted by Koo & McKee (1992) for a wind expanding into a r^{-2} density profile. The symmetry of the winds means that instability is seeded by the integration errors associated with the grid discretization, and so the solution is artificially symmetric compared with a real nebula. Implementation of some random or clumpy component to the wind boundary condition would break this symmetry and produce more realistic nebulae. We plan higher resolution simulations of this wind-wind interaction for comparison with observations of WR nebulae.

Finally, we revisited the 2D simulation by Stevens et al. (1992) of the wind-wind collision in the WR + O-star binary system V444 Cyg, using 3D MHD simulations including moderate stellar rotation. We find very similar results for the hydrodynamics of this system, which is marginally in the regime where strong cooling is expected to produce thin-shell instabilities and strong distortions of the bow shock structure. This proof-of-concept calculation requires a number of enhancements in order to reach state of the art, especially the inclusion of orbital motion and a better implementation of radiative cooling for hydrogen-poor plasmas. It is nevertheless encouraging that one can obtain good results with modest computational resources, and we intend to develop this setup significantly in future work.

The parallel scaling of PION is shown to be very good for MHD calculations without radiative transfer, as long as each MPI process has a local subdomain of $\gtrsim 32^3$ grid cells per level. Beyond this, the ratio of boundary cells to grid cells becomes sufficiently large that the computation and communication overhead is prohibitive. We anticipate that this scaling can be further improved significantly by implementation of hybrid MPI + OpenMP parallelization, because of reduced communication overhead and fewer boundary cells with duplicated computation. Scaling for 3D R-HD simulations is less good, losing two times in efficiency when the number of MPI processes increases by eight times, and significant work is needed to get this running efficiently on large supercomputers.

We introduced the PYPION library of PYTHON routines for reading PION snapshots and making various plots of the gas properties as a function of position. We demonstrate a method for converting PION snapshots to FITS images that can be read by the TORUS Monte Carlo radiative transfer code (Harries et al. 2019) and post-processed to calculate thermal dust emission maps. This method can also be extended to enable plotting of emission maps from spectral lines, thermal X-rays, or any form of radiation where the emissivity is a simple function of density and temperature.

It is hoped that PION will be a useful tool for the community to model nebulae around evolving massive stars. The source code can be obtained via a git repository from <https://www.pion.ie/>, and contributed code can be added using a mirrored repository on GitHub. A mailing list is also available for user support and discussion at <https://groups.io/g/pion>. The methods developed here will be used in future work to study bow shocks and wind-blown nebulae around massive stars with 3D simulations. Comparing synthetic and real observations will allow us to test the predictions of stellar evolution calculations, to learn more about stellar mass-loss, magnetism, and particle acceleration.

ACKNOWLEDGEMENTS

JM is grateful to N. Langer for advice and discussions on circumstellar medium modelling, and for providing evolutionary calculations used in Sections 4.2 and 4.3, to Jim Stone and Robin Williams for advice and suggestions on mesh refinement, and to Luca Grassitelli for useful discussions on winds from evolved massive stars. JM acknowledges funding from a Royal Society-Science Foundation Ireland University Research Fellowship (14/RS-URF/3219). DZ acknowledges funding from an Irish Research Council (IRC) Starting Laureate Award (IRCLA\2017\83). SG is funded by a Hamilton Scholarship from the Dublin Institute for Advanced Studies. MM acknowledges funding from a Royal Society Research Fellows Enhancement Award (RGF\EA\180214). TJH is funded by a Royal Society Dorothy Hodgkin Fellowship. The authors wish to acknowledge the DJEI/DES/SFI/HEA Irish Centre for High-End Computing (ICHEC) for the provision of computational facilities and support (projects dsast022c, dsast023b). This research has made use of NASA's Astrophysics Data System Bibliographic Services, the GNU Scientific Library (Galassi et al. 2019) and, for visualization and analysis, VISIT (Childs et al. 2012), NUMPY (Harris et al. 2020), and MATPLOTLIB (Hunter 2007). This research made use of ASTROPY, a community-developed core PYTHON package for Astronomy (Astropy Collaboration et al. 2013, 2018).

DATA AVAILABILITY STATEMENT

Instructions for downloading source code for PION and PYPION can be found at <https://www.pion.ie/>. Source code for TORUS can be obtained

from <https://bitbucket.org/tjharries/torus/>. The data underlying this article will be shared on reasonable request to the corresponding author.

REFERENCES

- Arthur S. J., Henney W. J., Mellema G., de Colle F., Vázquez-Semadeni E., 2011, *MNRAS*, 414, 1747
- Astropy Collaboration et al., 2013, *A&A*, 558, A33
- Astropy Collaboration et al., 2018, *AJ*, 156, 123
- Beasor E. R., Davies B., Smith N., van Loon J. T., Gehr R. D., Figer D. F., 2020, *MNRAS*, 492, 5994
- Benaglia P., Romero G. E., Martí J., Peri C. S., Araudo A. T., 2010, *A&A*, 517, L10
- Berger M. J., Colella P., 1989, *J. Comput. Phys.*, 82, 64 (BC89)
- Bestenlehner J. M., 2020, *MNRAS*, 493, 3938
- Bisbas T. G. et al., 2015, *MNRAS*, 453, 1324
- Bjorkman J. E., Cassinelli J. P., 1993, *ApJ*, 409, 429
- Bodenheimer P., Tenorio-Tagle G., Yorke H. W., 1979, *ApJ*, 233, 85
- Cantiello M. et al., 2009, *A&A*, 499, 279
- Cargo P., Gallice G., 1997, *J. Comput. Phys.*, 136, 446
- Childs H. et al., 2012, Lawrence Berkeley National Laboratory series, High Performance Visualization-Enabling Extreme-Scale Scientific Insight. eScholarship Publishing, California, USA, p. 357
- Chita S. M., Langer N., van Marle A. J., García-Segura G., Heger A., 2008, *A&A*, 488, L37 (CVL08)
- Dale J. E., 2015, *New A Rev.*, 68, 1
- Daley-Yates S., Stevens I. R., ud-Doula A., 2019, *MNRAS*, 489, 3251
- Dedner A., Kemm F., Kröner D., Munz C.-D., Schnitzer T., Wesenberg M., 2002, *J. Comput. Phys.*, 175, 645
- del Valle M. V., Pohl M., 2018, *ApJ*, 864, 19
- Derigs D., Winters A. R., Gassner G. J., Walch S., Böhm M., 2018, *J. Comput. Phys.*, 364, 420
- Draine B. T., 2003, *ARA&A*, 41, 241
- Eldridge J. J., Genet F., Daigne F., Mochkovitch R., 2006, *MNRAS*, 367, 186
- Falle S. A. E. G., 1991, *MNRAS*, 250, 581
- Falle S., Komissarov S., Joarder P., 1998, *MNRAS*, 297, 265
- Fossati L. et al., 2015, *A&A*, 582, A45
- Frank A., Mellema G., 1994, *A&A*, 289, 937
- Freyer T., Hensler G., Yorke H. W., 2003, *ApJ*, 594, 888
- Freyer T., Hensler G., Yorke H. W., 2006, *ApJ*, 638, 262
- Galassi M. et al., 2019, GNU Scientific Library Reference Manual, 3rd edn. Network Theory Ltd., United Kingdom. Available at: <http://www.gnu.org/software/gsl/>
- García-Segura G., Mac Low M., Langer N., 1996a, *A&A*, 305, 229
- García-Segura G., Langer N., Mac Low M., 1996b, *A&A*, 316, 133
- Gardiner T. A., Stone J. M., 2005, *J. Comput. Phys.*, 205, 509
- Geen S., Rosdahl J., Blaizot J., Devriendt J., Slyz A., 2015a, *MNRAS*, 448, 3248
- Geen S., Hennebelle P., Tremblin P., Rosdahl J., 2015b, *MNRAS*, 454, 4484
- Georgy C., Walder R., Folini D., Bykov A., Marcowith A., Favre J. M., 2013, *A&A*, 559, A69
- Grassitelli L., Fossati L., Langer N., Miglio A., Istrate A. G., Sanyal D., 2015, *A&A*, 584, L2
- Grassitelli L., Langer N., Grin N. J., Mackey J., Bestenlehner J. M., Gräfenber G., 2018, *A&A*, 614, A86
- Grassitelli L., Langer N., Mackey J., Graefener G., Grin N., Sander A., Vink J., 2021, *A&A*, 647, 14
- Green S., Mackey J., Haworth T. J., Gvaramadze V. V., Duffy P., 2019, *A&A*, 625, A4
- Grimaldo E., Reimer A., Kissmann R., Niederwanger F., Reitberger K., 2019, *ApJ*, 871, 55
- Gvaramadze V. V. et al., 2015, *MNRAS*, 454, 219
- Gvaramadze V. V., Mackey J., Kniazev A. Y., Langer N., Chené A. N., Castro N., Haworth T. J., Grebel E. K., 2017, *MNRAS*, 466, 1857
- Haid S., Walch S., Seifried D., Wunsch R., Dinnbier F., Naab T., 2018, *MNRAS*, 478, 4799

- Harries T. J., Haworth T. J., Acreman D., Ali A., Douglas T., 2019, *Astron. Comput.*, 27, 63
- Harris C. R. et al., 2020, *Nature*, 585, 357
- Harten A., Lax P. D., Leer B. v., 1983, *SIAM Rev.*, 25, 35
- Heger A., Langer N., 2000, *ApJ*, 544, 1016
- Henney W. J., Arthur S. J., de Colle F., Mellema G., 2009, *MNRAS*, 398, 157
- Hummer D. G., 1994, *MNRAS*, 268, 109
- Humphreys R. M., Helmelt G., Jones T. J., Gordon M. S., 2020, *AJ*, 160, 145
- Hunter J. D., 2007, *Comput. Sci. & Eng.*, 9, 90
- Innes D., Giddings J., Falle S., 1987, *MNRAS*, 226, 67
- Janhunen P., 2000, *J. Comput. Phys.*, 160, 649
- Keppens R., Teunissen J., Xia C., Porth O., 2021, *Computers & Mathematics with Applications*, 81, 316
- Kissmann R., Reitberger K., Reimer O., Reimer A., Grimaldo E., 2016, *ApJ*, 831, 121
- Koo B.-C., McKee C. F., 1992, *ApJ*, 388, 103
- Lamberts A., Fromang S., Dubus G., 2011, *MNRAS*, 418, 2618
- Lamberts A. et al., 2017, *MNRAS*, 468, 2655
- Langer N., 1997, in Nota A., Lamers H., eds, *ASP Conf. Ser. Vol. 120, Luminous Blue Variables: Massive Stars in Transition*. Astron. Soc. Pac., San Francisco, p. 83
- Langer N., 2012, *ARA&A*, 50, 107
- Langer N., García-Segura G., Mac Low M., 1999, *ApJ*, 520, L49
- Mac Low M.-M., McCray R., Norman M. L., 1989, *ApJ*, 337, 141
- Mac Low M.-M., van Buren D., Wood D. O. S., Churchwell E., 1991, *ApJ*, 369, 395
- Mackey J., 2012, *A&A*, 539, A147
- Mackey J., Lim A. J., 2010, *MNRAS*, 403, 714
- Mackey J., Lim A. J., 2011, *MNRAS*, 412, 2079
- Mackey J., Mohamed S., Neilson H. R., Langer N., Meyer D. M.-A., 2012, *ApJ*, 751, L10
- Mackey J., Langer N., Gvaramadze V. V., 2013, *MNRAS*, 436, 859
- Mackey J., Mohamed S., Gvaramadze V. V., Kotak R., Langer N., Meyer D. M.-A., Moriya T. J., Neilson H. R., 2014, *Nature*, 512, 282
- Mackey J., Gvaramadze V. V., Mohamed S., Langer N., 2015, *A&A*, 573, A10
- Mackey J., Haworth T. J., Gvaramadze V. V., Mohamed S., Langer N., Harries T. J., 2016, *A&A*, 586, A114
- Mackey J., Green S., Moutzouri M., 2020, *J. Phys.*, 1620, 012012
- Madura T. I. et al., 2013, *MNRAS*, 436, 3820
- Mathis J. S., Rimpl W., Nordsieck K. H., 1977, *ApJ*, 217, 425
- Meliani Z., Keppens R., Casse F., Giannios D., 2007, *MNRAS*, 376, 1189
- Mellema G., 1994, *A&A*, 290, 915
- Mellema G., Iliev I., Alvarez M., Shapiro P., 2006a, *New Astron.*, 11, 374
- Mellema G., Arthur S., Henney W., Iliev I., Shapiro P., 2006b, *ApJ*, 647, 397
- Meyer D. M.-A., Mackey J., Langer N., Gvaramadze V. V., Mignone A., Izzard R. G., Kaper L., 2014, *MNRAS*, 444, 2754
- Meyer D. M.-A., Langer N., Mackey J., Velázquez P. F., Gusdorf A., 2015, *MNRAS*, 450, 3080
- Meyer D. M.-A., Mignone A., Kuiper R., Raga A. C., Kley W., 2017, *MNRAS*, 464, 3229
- Meyer D. M. A., Petrov M., Pohl M., 2020, *MNRAS*, 493, 3548
- Mignone A., Zanni C., Tzeferacos P., van Straalen B., Colella P., Bodo G., 2012, *ApJS*, 198, 7
- Miyoshi T., Kusano K., 2005, *J. Comput. Phys.*, 208, 315
- Mohamed S., Mackey J., Langer N., 2012, *A&A*, 541, A1
- Parker E. N., 1958, *ApJ*, 128, 664
- Parkin E. R., Gosset E., 2011, *A&A*, 530, A119
- Parkin E. R., Pittard J. M., Corcoran M. F., Hamaguchi K., 2011, *ApJ*, 726, 105
- Peri C. S., Benaglia P., Brookes D. P., Stevens I. R., Isequilla N. L., 2012, *A&A*, 538, A108
- Peri C. S., Benaglia P., Isequilla N. L., 2015, *A&A*, 578, A45
- Pittard J. M., 2009, *MNRAS*, 396, 1743
- Pittard J. M., Dougherty S. M., 2006, *MNRAS*, 372, 801
- Pittard J. M., Parkin E. R., 2010, *MNRAS*, 403, 1657
- Plewa T., Müller E., 1999, *A&A*, 342, 179
- Pogorelov N. V., Zank G. P., Ogino T., 2004, *ApJ*, 614, 1007
- Pogorelov N. V., Zank G. P., Ogino T., 2006, *ApJ*, 644, 1299
- Pogorelov N. V., Suess S. T., Borovikov S. N., Ebert R. W., McComas D. J., Zank G. P., 2013, *ApJ*, 772, 2
- Powell K., Roe P., Linde T., Gombosi T., de Zeeuw D., 1999, *J. Comput. Phys.*, 154, 284
- Puls J., Vink J. S., Najarro F., 2008, *A&A Rev.*, 16, 209
- Raga A. C., Noriega-Crespo A., Cantó J., Steffen W., van Buren D., Mellema G., Lundqvist P., 1997, *Rev. Mex. Astron. Astrofis.*, 33, 73
- Rogers H., Pittard J. M., 2013, *MNRAS*, 431, 1337
- Rosen A. L., Krumholz M. R., Oishi J. S., Lee A. T., Klein R. I., 2017, *J. Comput. Phys.*, 330, 924
- Rozyczka M., 1985, *A&A*, 143, 59
- Rybicki G. B., Lightman A. P., 1979, *Radiative Processes in Astrophysics*. John Wiley & Sons, Inc., New York
- Sander A. A. C., Vink J. S., 2020, *MNRAS*, 499, 873
- Scherer K., Baalmann L. R., Fichtner H., Kleimann J., Bomans D. J., Weis K., Ferreira S. E. S., Herbst K., 2020, *MNRAS*, 493, 4172
- Smith N., 2014, *ARA&A*, 52, 487
- Smith N. et al., 2018, *MNRAS*, 480, 1466
- Snow T. P., Jr, Morton D. C., 1976, *ApJS*, 32, 429
- Stevens I. R., Blondin J. M., Pollock A. M. T., 1992, *ApJ*, 386, 265
- Stone J., Gardiner T., Teuben P., Hawley J., Simon J., 2008, *ApJS*, 178, 137
- Stone J. M., Tomida K., White C. J., Felker K. G., 2020, *ApJS*, 249, 4
- Szécsi D., Mackey J., Langer N., 2018, *A&A*, 612, A55
- Tóth G., Roe P. L., 2002, *J. Comput. Phys.*, 180, 736
- van Marle A. J., Keppens R., 2012, *A&A*, 547, A3
- van Marle A. J., Langer N., García-Segura G., 2005, *A&A*, 444, 837
- van Marle A. J., Langer N., Yoon S., García-Segura G., 2008, *A&A*, 478, 769
- van Marle A. J., Decin L., Meliani Z., 2014, *A&A*, 561, A152
- van Veelen B., Langer N., Vink J., García-Segura G., van Marle A. J., 2009, *A&A*, 503, 495
- Walch S. K., Whitworth A. P., Bisbas T., Wunsch R., Hubber D., 2012, *MNRAS*, 427, 625
- Walder R., Folini D., Meynet G., 2012, *Space Sci. Rev.*, 166, 145
- Washimi H., Tanaka T., 2001, *Adv. Space Res.*, 27, 509
- White R., Breuhaus M., Konno R., Ohm S., Reville B., Hinton J. A., 2020, *A&A*, 635, A144
- Wiersma R. P. C., Schaye J., Smith B. D., 2009, *MNRAS*, 393, 99
- Williams R. J. R., Bisbas T. G., Haworth T. J., Mackey J., 2018, *MNRAS*, 479, 2016
- Wise J. H., Abel T., 2011, *MNRAS*, 414, 3458
- Yorke H. W., Kaisig M., 1995, *Comput. Phys. Commun.*, 89, 29
- Yorke H. W., Welz A., 1996, *A&A*, 315, 555

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.