

OCTO-TIGER: a new, 3D hydrodynamic code for stellar mergers that uses HPX parallelization

Dominic C. Marcelllo,^{1,2★†} Sagiv Shiber,^{1★†} Orsola De Marco,^{3,4} Juhan Frank,¹ Geoffrey C. Clayton¹,¹ Patrick M. Motl,⁵ Patrick Diehl^{2★} and Hartmut Kaiser²

¹Department of Physics and Astronomy, Louisiana State University, Baton Rouge, LA 70803 USA

²Center for Computational Technologies, Louisiana State University, Baton Rouge, LA 70803 USA

³Department of Physics and Astronomy, Macquarie University, Sydney, NSW 2109, Australia

⁴Astronomy, Astrophysics and Astrophotonics Research Centre, Macquarie University, Sydney, NSW 2109, Australia

⁵The School of Sciences, Indiana University Kokomo, Kokomo, IN 46904, USA

Accepted 2021 March 22. Received 2021 March 22; in original form 2020 December 25

ABSTRACT

OCTO-TIGER is an astrophysics code to simulate the evolution of self-gravitating and rotating systems of arbitrary geometry based on the fast multipole method, using adaptive mesh refinement. OCTO-TIGER is currently optimized to simulate the merger of well-resolved stars that can be approximated by barotropic structures, such as white dwarfs (WDs) or main-sequence stars. The gravity solver conserves angular momentum to machine precision, thanks to a ‘correction’ algorithm. This code uses HPX parallelization, allowing the overlap of work and communication and leading to excellent scaling properties, allowing for the computation of large problems in reasonable wall-clock times. In this paper, we investigate the code performance and precision by running benchmarking tests. These include simple problems, such as the Sod shock tube, as well as sophisticated, full, WD binary simulations. Results are compared to analytical solutions, when known, and to other grid-based codes such as FLASH. We also compute the interaction between two WDs from the early mass transfer through to the merger and compare with past simulations of similar systems. We measure OCTO-TIGER’s scaling properties up to a core count of $\sim 80\,000$, showing excellent performance for large problems. Finally, we outline the current and planned areas of development aimed at tackling a number of physical phenomena connected to observations of transients.

Key words: hydrodynamics – methods: analytical – methods: numerical – binaries: close – stars: evolution – white dwarfs.

1 INTRODUCTION

Approximately two-thirds of observed stars are members of binary or multiple stellar systems (Duquennoy & Mayor 1991). While the evolution of single stars is predictable given the initial mass and composition, the evolution of the components of a binary or multiple system can be dramatically altered if they are close enough to become interactive through mass transfer and mutual irradiation. The most consequential interactions in binary systems occur when the components are so close that one (or both) of the stars fill their Roche lobes and mass is transferred from one star to the other. In some cases, the transfer proceeds steadily for a long time; in others mass and energy are exchanged in a common envelope; while in the most dynamic and interesting cases, the interaction leads to a merger giving birth to a single star of unusual properties that cannot be produced by the evolution of a single star. We know that all of these processes do occur in nature yielding observable outcomes such as contact binaries (Smith 1984; Rucinski 2010), helium white dwarfs (WD; Kippenhahn, Kohl & Weigert 1967; Kippenhahn, Thomas & Weigert

1968), R Coronae Borealis stars (Clayton 2012), Type Iax supernovae (Solheim 2010; Jha 2017), and dwarf novae (Frank, King & Raine 2002; Warner 2003). At least a fraction of all Type Ia supernovae, which are some of the most energetic explosion in the Universe, are likely caused by the Roche lobe-overflow induced merger of two WDs (Hillebrandt & Niemeyer 2000). This particular type of supernova is used as a ‘standard candle’ for measuring distances to other galaxies. Because of this, understanding Type Ia supernovae (SN) is important for understanding the origins of the Universe and its eventual fate.

Common envelope binary interactions (Ivanova et al. 2013) play a crucial role in the formation of close binaries of all types including X-ray binaries (Tauris & van den Heuvel 2014), cataclysmic variables (Meyer & Meyer-Hofmeister 1979; Webbink 1992), close binaries in planetary nebulae (De Marco et al. 2015), and possibly even massive stellar double black hole binaries (Ricker et al. 2019). Clearly, common envelope interactions are close cousins of mergers, the main differences being the final outcomes, which in turn depend on the types of stars involved, their mass ratios, and the energetics of mass-loss through unbinding and ejection of the envelope.

Important observational developments in the past decade or so have lent further credence to the theoretical ideas involving binary mergers and given further impetus to the development of accurate and efficient numerical tools for investigating binary mergers and

* E-mail: dmarce504@gmail.com (DCM); shibersag@gmail.com (SS); pdiehl@cct.lsu.edu (PD)

† The first two authors made equal contributions to this paper.

common envelope evolution. For example, in 2008 September, the contact binary, V1309 Sco, merged to form a luminous red nova (Tylenda et al. 2011). When the merger occurred, the system increased in brightness by a factor of approximately 10^4 . Mason et al. (2010) observed the outburst spectroscopically, confirming it as a luminous red nova. Because the Optical Gravitational Lensing Experiment (OGLE) observed V1309 Sco prior to its merger for 6 yr, this provided the rare opportunity to observe a stellar merger both before and after the event, prompting Tylenda to refer to V1309 Sco as the ‘Rosetta Stone’ of contact-binary mergers. More recently, the Zwicky Transient Factory discovery of an eclipsing double WD binary with an orbital period of 8.8 min, which is destined to merge and become a hot subdwarf or an R Coronae Borealis star, makes it clear that such progenitor binaries do exist and provides further motivation for investigating stellar mergers and their outcomes (Burdge et al. 2020). Discoveries of close-binary systems have multiplied in number due to surveys such as the Zwicky Transient Factory, and are due to increase exponentially with the arrival of the Vera Rubin Observatory in 2022 (Ivezić et al. 2007).

Ultimately, one would want a high-resolution 3D magneto- and radiation-hydrodynamic simulation with nuclear energy production, full chemistry, and resolution down to the dynamical time-scale while also being able to follow the system over thermal time-scales, but those capabilities are well beyond the horizon at this time. In lieu of such simulations, studies have used all manner of hybrid approaches. For example, 3D smooth particle hydrodynamics (SPH) simulations of merging WD binaries with 1.8 million particles are then mapped into grid codes to model SN detonations (Pakmor et al. 2012) and the output is then used to study the light and polarization signals (Bulla et al. 2016). In a different simulation, two core hydrogen burning massive stars are merged using the moving mesh code, AREPO, with a resolution of 400 000 to 4 million cells and the merged stars were then mapped into a 1D code to study the evolution of the remnant (Schneider et al. 2019). A simulation of the V1309 Sco merger was carried out using 100 000 SPH particles (Nandez, Ivanova & Lombardi 2014), and then used as a starting point for a detailed discussion aided by analytical physics. A number of 3D hydrodynamic simulations were also used to model other binary interactions such as the mass transfer preceding coalescence (e.g. Pejcha, Metzger & Tomida 2016; MacLeod & Loeb 2020). Also, Kashyap et al. (2018) carried out 200 000 particle SPH simulations of WD mergers that then were mapped into a grid code to investigate the detonation properties.

Much of the research in 3D hydrodynamic simulations of WD mergers is aimed at understanding the dynamics of detonation with a goal to understand Type Ia supernovae. In their detailed review of 3D, WD merger simulations, Katz et al. (2016) described several efforts, most of which concentrate on the details of the properties at the time of merger but do not necessarily model the early mass transfer and merger phase. This review cites the early pioneering efforts to model the entire merger by Motl, Tohline & Frank (2002), D’Souza et al. (2006), and Motl et al. (2007). In a recent counterpart to those papers, Motl et al. (2017) carried out simulations of merging WD stars, using a finite difference technique code with up to 4 million cells and an SPH code with up to 1 million particles. Aside from the importance of code comparison, that publication explicitly shows the resolution and wall-clock time constraints of this type of simulations.

In this paper, we present OCTO-TIGER, a code that aims at improving 3D simulations of interactions using a number of computational techniques that increase accuracy and scalability. This will allow us to calculate full mergers with reasonably high resolution, and reasonable wall-clock times and has the capacity to include a greater

amount of physics without moving the computation into the realm of impossibility. OCTO-TIGER also conserves energy and angular momentum to excellent precision. OCTO-TIGER’s main application currently is simulating the merger of well-resolved stars that can be represented via polytropes, such as main-sequence stars or WDs. Marcelllo et al. (2016) presented OCTO-TIGER, with a description of the governing equations along with some preliminary tests. Staff et al. (2018) and Kadam et al. (2018) used OCTO-TIGER in parallel with a suite of other codes to simulate WD binary mergers leading to R Coronae Borealis stars and contact binaries. These studies provided a test of sorts for OCTO-TIGER, but the scope of those papers was such that a systematic verification and validation of OCTO-TIGER was not carried out, and neither was a scaling test aimed at measuring its speed. In the meantime, several improvements have been implemented in the code. It is therefore appropriate and timely to test and document OCTO-TIGER in a systematic way, by carrying out a suite of standard benchmark simulations, a comparison to other codes, complete with scaling tests of the latest code version.

This paper is structured in the following way. In Section 2, we describe OCTO-TIGER’s underlying equations. In Section 3, we present the benchmarks, starting with the shock tube (Section 3.2), the Sedov blastwave (Section 3.3), a uniform static sphere (Section 3.4) and continuing with a static pulsating polytrope (Section 3.5), a translating polytrope (Section 3.6), and a rotating polytrope (Section 3.7). We conclude with a binary simulation with a mass ratio of 0.5 in Section 4, and an assessment of the scaling properties of OCTO-TIGER in Section 5. We conclude in Section 6.

2 THE AMR CODE OCTO-TIGER

OCTO-TIGER is an Eulerian adaptive mesh refinement (AMR) code, optimized for the simulation of inviscid, compressible fluids. Below we fully describe its governing equations and numerical methods.

2.1 Hydrodynamic evolution equations

OCTO-TIGER evolves Euler’s inviscid equations of motion for a self-gravitating fluid on a rotating mesh. The evolution equations are:

$$\frac{\partial}{\partial t} \rho_m + \nabla \cdot \rho_m \mathbf{v} = 0, \quad (1)$$

$$\frac{\partial}{\partial t} \mathbf{s} + \nabla \cdot \mathbf{v} \mathbf{s} + \nabla p = \rho \mathbf{g} - \Omega \times \mathbf{s}, \quad (2)$$

$$\begin{aligned} & \frac{\partial}{\partial t} \left(E + \frac{1}{2} \rho \phi \right) + \nabla \cdot \mathbf{v} (E + \rho \phi) + \nabla \cdot \mathbf{u} p \\ & = \frac{1}{2} \left(\phi \frac{\partial}{\partial t} \rho - \rho \frac{\partial}{\partial t} \phi \right) + \rho \Omega \cdot (\mathbf{x} \times \mathbf{g}), \end{aligned} \quad (3)$$

and

$$\nabla^2 \phi = 4\pi G \rho, \quad (4)$$

where ρ_m is the mass density of the m th species, \mathbf{v} is the velocity in the rotating frame, \mathbf{s} is the inertial frame momentum density, p is the gas pressure, \mathbf{u} is the inertial frame velocity, \mathbf{g} is the gravitational acceleration, Ω is the rotational frequency of the grid, E is the gas internal plus bulk kinetic energy density in the inertial frame, $\rho = \sum \rho_m$ is the total mass density, \mathbf{x} is the position vector on the grid, ϕ is the gravitational potential, and G is the gravitational constant. The rotating frame velocity is related to the inertial frame momentum density by

$$\mathbf{s} = \rho (\mathbf{v} + \Omega \times \mathbf{x}). \quad (5)$$

The gravitational potential is related to the gravitational acceleration by $\mathbf{g} = -\phi$. The time derivatives are taken in the rotating frame. These are related to time derivatives in the inertial frame by

$$\frac{\partial}{\partial t} = \left(\frac{\partial}{\partial t} \right)_{\text{inertial}} + (\mathbf{x} \times \boldsymbol{\Omega}) \cdot \nabla. \quad (6)$$

A previous version of OCTO-TIGER used a hydro-solver that conserved linear and angular momentum to machine precision. This was based on the work by Després & Labourasse (2015). We successfully adapted this method to OCTO-TIGER for use when velocities are reconstructed in the inertial frame. However, our methodologies could not be applied to when velocities are reconstructed in the rotating frame. The rotating frame is a better choice when dealing with binary interactions, because equilibrium rotating stars retain their initial profiles much better. We thus abandoned the development of the hydrodynamics angular momentum conservation feature. A legacy of this development work is a separately evolved angular momentum field, \mathbf{l} , described by the following equation:

$$\frac{\partial}{\partial t} \mathbf{l} + \nabla \cdot \mathbf{v} \mathbf{l} + \nabla \times \mathbf{x} p = \rho \mathbf{x} \times \mathbf{g} - \boldsymbol{\Omega} \times \mathbf{l}. \quad (7)$$

In the current version, this quantity is evolved passively, meaning the evolution of the other variables does not depend on the value of \mathbf{l} . It is initialized to $\mathbf{x} \times \mathbf{s}$ at $t = 0$. Although our gravity solver exactly conserves angular momentum to machine precision, because of numerical viscosity in the hydrodynamics solver, angular momentum is not exactly conserved. The difference between $\mathbf{x} \times \mathbf{s}$ and \mathbf{l} after $t = 0$ gives a measure of the error in angular momentum conservation.

The source terms on the right-hand side (RHS) of equation (2) include a gravitational and a rotational term. The rotational term accounts for the rotation of the momentum vector relative to the rotating mesh. Because we evolve inertial frame quantities on a rotating mesh, as opposed to rotating frame quantities on a rotating mesh, this term is half the Coriolis force (only half because velocities are taken with respect to the rotating grid, but momenta are calculated in the inertial frame). Note that if $\int_V \mathbf{s} = 0$, this term does not violate momentum conservation.

Rather than solving equation (4) using an iterative approach, OCTO-TIGER uses the fast multipole method (FMM; described below). To calculate the first term on the RHS of equation (3), we solve for $\frac{\partial}{\partial t} \phi$ using the FMM with the numerically computed value of $\frac{\partial}{\partial t} \rho$ as the source term. This results in the two parts of this term cancelling when summed over the entire grid.

Equation (3) is derived from the usual form of the energy equation:

$$\frac{\partial}{\partial t} E + \nabla \cdot \mathbf{v} E - \nabla \cdot \mathbf{u} p = -\rho \mathbf{u} \cdot \nabla \phi, \quad (8)$$

along with equation (1). Equation (3) is written in a form that emphasizes that the conserved quantity is the gas energy, E (internal plus bulk kinetic), plus the potential energy or $E + \frac{1}{2} \rho \phi$. The first term on the RHS of equation (3) vanishes globally because ϕ is linearly related to ρ , while the second term vanishes globally because the total change in angular momentum due to gravity over all space is zero. As discussed by Marcelllo & Tohline (2012), evolving the energy equation in this form prevents violation of energy conservation due to matter moving up or down a potential well because of numerical viscosity. Because our gravity solver conserves linear and angular momenta to machine precision, using this form of the energy equation conserves total energy to machine precision. In practice, we actually evolve E instead of $E + \frac{1}{2} \rho \phi$, by taking the discretized evolution equation for equation (3) and solving it for E .

Following the dual energy formalism of Bryan et al. (1995), OCTO-TIGER evolves a second variable for the energy. While they used the internal energy density as the second energy variable, we use the ‘entropy tracer’ (Motl et al. 2002), defined as

$$\tau = (\rho \varepsilon)^{\frac{1}{\gamma}}, \quad (9)$$

where ε is the specific internal gas energy and $\gamma = \frac{5}{3}$ is the ratio of specific heats. When there are no shocks, τ is conserved and evolves as

$$\frac{\partial}{\partial t} \tau + \nabla \cdot \tau \mathbf{v} = 0. \quad (10)$$

The specific internal energy, ε , is computed according to

$$\rho \varepsilon = \begin{cases} E - \frac{1}{2} \rho u^2, & \text{if } E - \frac{1}{2} \rho u^2 \geq \epsilon_1 E \\ \tau^\gamma, & \text{otherwise} \end{cases}, \quad (11)$$

where the default value of $\epsilon_1 = 0.001$. Once we have the internal energy density, we can compute the pressure with the ideal gas equation:

$$p = (\gamma - 1) \rho \varepsilon. \quad (12)$$

At the end of every update of the evolution variables, the entropy tracer is then reset using equation (9) in computational cells, which satisfy

$$E - \frac{1}{2} \rho u^2 > \epsilon_2 E \quad (13)$$

for at least one of the adjacent cells or the cell itself, where the default value of $\epsilon_2 = 0.1$. Otherwise it is left alone. We use the values for ϵ_1 and ϵ_2 chosen by Bryan et al. (1995).

This treatment allows for the proper evolution of shocks while simultaneously retaining numerical accuracy of the internal energy in high mach flows, where the kinetic energy dominates. It is roughly analogous to adding extra digits of precision to the internal energy. This method also guarantees positive values for the internal energy even when $E - \frac{1}{2} \rho u^2 < 0$.

OCTO-TIGER evolves the gas using an ideal gas equation of state (EoS), but by setting $\epsilon_1 = \epsilon_2 = 1$, shock heating is eliminated and effectively OCTO-TIGER evolves the gas with a polytropic EoS. If

$$p_{\text{poly}} = K \rho^{1+\frac{1}{n}}, \quad (14)$$

initially, where K and n are constants, this condition will continue to hold as the system evolves.

2.2 Gravity update and the angular momentum correction

OCTO-TIGER uses a variant of the Cartesian FMM described by Dehnen (2000). This method conserves linear momentum to machine precision. The FMM adapts naturally to an oct-tree based AMR scheme such as that used by OCTO-TIGER. The multipoles of each cell are composed of the multipoles of its child cells, and the expansion in each child cell is derived from the expansion of its parent cell. Both multipole and expansions are relative to cell centres.

OCTO-TIGER has an extension, detailed by Marcelllo et al. (2016), which allows its gravity solver to conserve both angular and linear momentum to machine precision. Based on the FMM described by Dehnen (2000), which conserves linear momentum to machine precision, the OCTO-TIGER extension works by adding an additional higher order multipoles that is used to calculate a correction to the calculated force. This angular momentum correction cancels the angular momentum conservation violation while preserving linear momentum conservation of the original method. As a result of angular momentum conservation, the last term in the RHS of

equation (3) also sums to zero over all space, enabling conservation of energy to machine precision in the rotating frame.

Conservation of angular momentum and energy is important for the self-consistent accurate modelling of near equilibrium astrophysical systems such as binary stars in the early phases of mass transfer. Violations in energy conservation can cause stars in equilibrium to ‘evaporate’, while the balance of angular momentum may affect the ultimate merger if one occurs.

The FMM requires the specification of an opening criterion. Given two cells, A and B , at grid locations \mathbf{x}_A and \mathbf{x}_B , respectively, the opening angle is

$$\theta' = \frac{\Delta x}{|\mathbf{x}_A - \mathbf{x}_B|}, \quad (15)$$

where Δx is the width of a grid cell. If $\theta' < \theta$, where θ is a critical value less than unity, then two cells are well separated. Two well separated cells interact through the multipole interaction when they are well separated from each other, but their respective parents are not well separated. In the models presented in this paper, $\theta = 0.5$ or 0.34 , the latter being used for the interacting binary models.

We also define the quantity, θ_{\min} , to refer to the lowest θ allowed by OCTO-TIGER. This is determined at compile time and its value depends on the size of the subgrids, with larger subgrids allowing for smaller values of θ_{\min} . For $8 \times 8 \times 8$ subgrids, as used in this paper, $\theta_{\min} = 0.34$.

It should be noted the current version of OCTO-TIGER does not conserve angular momentum in the hydrodynamics module. We have experimented with extending an angular momentum conserving hydrodynamics method described by Després & Labourasse (2015) for use in OCTO-TIGER. While we have had success when the grid is not rotating, the method fails in the rotating frame. The reconstruction of face values is computed in the rotating frame, eliminating the degree of freedom in the reconstruction required for Després & Labourasse (2015) to work.

2.3 Hydrodynamic update

The hydrodynamic update begins by reconstructing cell averaged values on the surface of the cells. OCTO-TIGER reconstructs values at the geometric centres of the 6 cell faces, 12 cell edges, and 8 cell vertices, for a total of 26 quadrature points. Note this results in each face having 9 quadrature points, 1 at the cell face centre, and 4 each for the cell face’s edges and vertices. This allows OCTO-TIGER to compute the flux through a face as the integral of the fluxes at 9 points on the face.

Rather than reconstructing the vector of conserved quantities,

$$U = \begin{bmatrix} \rho_m \\ \mathbf{s} \\ E \\ \tau \\ \rho\phi \end{bmatrix}, \quad (16)$$

we reconstruct

$$V = \begin{bmatrix} \rho_m \\ \mathbf{v} \\ E - \frac{1}{2}\rho u^2 \\ \tau \\ \phi \end{bmatrix}, \quad (17)$$

and then transform back to U . The reconstruction of velocities is done in the rotating frame. In OCTO-TIGER, the rotating frame is defined as rotating around the z -axis, and the frequency can be specified by

the user. For the binary simulations in this paper, the rotating frame frequency is the same as the initial orbital frequency of the binary. Reconstructing the velocities in the inertial frame is also a valid choice. However, we have found that single stars in initial equilibrium retain their original density profiles better when reconstructing the rotating frame velocities.

The total energy, defined as the sum of the potential and gas energy, is a conserved quantity, but it is not a purely hydrodynamic quantity. The local potential plays virtually no role in the formation of discontinuities in the gas energy. When applying a central advection scheme, we want the averaging of left-hand and right-hand states to occur over hydrodynamic variables. Furthermore, discontinuities in the potential energy are caused solely by the local mass density, the specific potential energy itself being smooth. For these reasons, we treat the two quantities as separate when computing the reconstruction and fluxes.

The left-hand side of equation (3) is evolved as the sum of two parts,

$$\frac{\partial}{\partial t} E + \nabla \cdot E\mathbf{v} + \nabla \cdot E\mathbf{u}p = 0, \quad (18)$$

and

$$\frac{\partial}{\partial t} \frac{1}{2}\rho\phi + \nabla \cdot \mathbf{v}\rho\phi = 0. \quad (19)$$

Below we show how we combine these two parts to form a single energy equation for the gas energy.

There are left and right values for each interface, and here we denote those with R and L superscripts, respectively. For a quantity u and for every i, j , and k there are eight vertex values,

$$\begin{aligned} &u_{i+1/2j+1/2k+1/2}^{RRR}, u_{i+1/2j+1/2k+1/2}^{RRL}, u_{i+1/2j+1/2k+1/2}^{RLR}, \\ &u_{i+1/2j+1/2k+1/2}^{LLL}, u_{i+1/2j+1/2k+1/2}^{LRR}, u_{i+1/2j+1/2k+1/2}^{LRL}, \\ &u_{i+1/2j+1/2k+1/2}^{LLR}, \text{ and } u_{i+1/2j+1/2k+1/2}^{LLL}, \end{aligned} \quad (20)$$

12 edge values,

$$\begin{aligned} &u_{i+1/2j+1/2k}^{RR}, u_{i+1/2j+1/2k}^{RL}, u_{i+1/2j+1/2k}^{LR}, u_{i+1/2j+1/2k}^{LL}, \\ &u_{i+1/2jk+1/2}^{RR}, u_{i+1/2jk+1/2}^{RL}, u_{i+1/2jk+1/2}^{LR}, u_{i+1/2jk+1/2}^{LL}, \\ &u_{ij+1/2k+1/2}^{RR}, u_{ij+1/2k+1/2}^{RL}, u_{ij+1/2k+1/2}^{LR}, \text{ and } u_{ij+1/2k+1/2}^{LL}, \end{aligned} \quad (21)$$

and six face values,

$$\begin{aligned} &u_{i+1/2jk}^R, u_{i+1/2jk}^L, \\ &u_{ij+1/2k}^R, u_{ij+1/2k}^L, \\ &u_{ijk+1/2}^R, \text{ and } u_{ijk+1/2}^L. \end{aligned} \quad (22)$$

Note that the u used here is not the inertial frame velocity used elsewhere in this paper. We use the piecewise parabolic method (PPM) of Colella & Woodward (1984, with contact discontinuity detection) to compute these values. The five cell stencil required for PPM is taken along the line from the cell centre to each face, edge, or vertex. For example, the reconstructed values at $u_{i+1/2jk}^L$ and $u_{i-1/2jk}^R$ are computed by applying PPM to the five cell stencil formed by $u_{i-2,j,k}$, $u_{i-1,j,k}$, $u_{i,j,k}$, $u_{i+1,j,k}$, and $u_{i+2,j,k}$, while the reconstructed values at $u_{i+1/2j-1/2k}^L$ and $u_{i-1/2j+1/2k}^R$ are computed with the stencil formed by $u_{i-2,j+2,k}$, $u_{i-1,j+1,k}$, $u_{i,j,k}$, $u_{i+1,j-1,k}$, and $u_{i+2,j-2,k}$. Reconstructing the primitive variables at 26 points across the cell’s surface, using PPM, results in multidimensional third-order convergence. Without this feature, the atmospheres of equilibrium stars turn into box like structures.

Once a cell’s evolved quantities are reconstructed at each of the 26 quadrature points on the cells’ surfaces, OCTO-TIGER computes

the fluxes at each of the 9 quadrature points for each cell's face. We use the central-upwind scheme described by Kurganov, Noelle & Petrova (2000). This scheme was originally chosen because it was straightforward to adapt it for use with OCTO-TIGER's hydrodynamics angular momentum conservation feature, as described in Section 2.1. It can be defined in terms of left and right face values,

$$H(U_L, U_R) = \frac{a^+ F(U_L) - a^- F(U_R)}{a^+ - a^-} + \frac{a^+ a^-}{a^+ - a^-} [U_R - U_L], \quad (23)$$

where H is the numerical flux, F is the physical flux, and a^+ and a^- are the positive and negative signal speeds. OCTO-TIGER obtains H for each dimension using the physical flux,

$$F_x = \begin{bmatrix} v_x \rho \\ v_x s_x + p \\ v_x s_y \\ v_x s_z \\ v_x E + u_x p \\ v_x \tau \\ v_x \rho \phi \end{bmatrix}, \quad F_y = \begin{bmatrix} v_y \rho \\ v_y s_x \\ v_y s_y + p \\ v_y s_z \\ v_y E + u_y p \\ v_y \tau \\ v_y \rho \phi \end{bmatrix}, \quad F_z = \begin{bmatrix} v_z \rho \\ v_z s_x \\ v_z s_y \\ v_z s_z + p \\ v_z E + u_z p \\ v_z \tau \\ v_z \rho \phi \end{bmatrix}, \quad (24)$$

and the signal speeds

$$\begin{aligned} a_x^+ &= \max(c_{s,R} + v_{x,R}, c_{s,L} + v_{x,L}, 0), \\ a_x^- &= \min(c_{s,R} - v_{x,R}, c_{s,L} - v_{x,L}, 0), \\ a_y^+ &= \max(c_{s,R} + v_{y,R}, c_{s,L} + v_{y,L}, 0), \\ a_y^- &= \min(c_{s,R} - v_{y,R}, c_{s,L} - v_{y,L}, 0), \\ a_z^+ &= \max(c_{s,R} + v_{z,R}, c_{s,L} + v_{z,L}, 0), \\ a_z^- &= \min(c_{s,R} - v_{z,R}, c_{s,L} - v_{z,L}, 0), \end{aligned} \quad (25)$$

where the c_s is the sound speed,

$$c_s = \sqrt{\frac{\gamma p}{\rho}}. \quad (26)$$

The total flux through a face is then obtained by summing the fluxes taken at each quadrature point on the face. For the x-face, for instance, we have

$$\begin{aligned} \mathcal{H}_{x,i+1/2jk} &= \frac{16}{36} [H_x(U_{i+1/2jk}^L, U_{i+1/2jk}^R)] \\ &+ \frac{4}{36} [H_x(U_{i+1/2j+1/2k}^{LL}, U_{i+1/2j+1/2k}^{RL})] \\ &+ H_x(U_{i+1/2j-1/2k}^{LR}, U_{i+1/2j-1/2k}^{RR}) \\ &+ H_x(U_{i+1/2j+1/2}^{LL}, U_{i-1/2jk+1/2}^{RL}) \\ &+ H_x(U_{i+1/2jk-1/2}^{LR}, U_{i-1/2jk-1/2}^{RR}) \\ &+ \frac{1}{36} [H_x(U_{i+1/2j+1/2k+1/2}^{LLL}, U_{i+1/2j+1/2k+1/2}^{RLL})] \\ &+ H_x(U_{i+1/2j+1/2k-1/2}^{LLR}, U_{i+1/2j+1/2k-1/2}^{RRL}) \\ &+ H_x(U_{i+1/2j-1/2k+1/2}^{LRL}, U_{i+1/2j-1/2k+1/2}^{RRL}) \\ &+ H_x(U_{i+1/2j-1/2k-1/2}^{LRR}, U_{i+1/2j-1/2k-1/2}^{RRR})], \end{aligned} \quad (27)$$

where $\mathcal{H}_{x,i+1/2jk}$ is the total numerical flux through the x-face.

The semidiscrete form of the evolution equations then becomes

$$\begin{aligned} \frac{d}{dt} U_{ijk} &+ \frac{H_{x,i+1/2jk} - H_{x,i-1/2jk}}{\Delta x} \\ &+ \frac{H_{y,ij+1/2k} - H_{y,ij-1/2k}}{\Delta x} \\ &+ \frac{H_{z,ijk+1/2} - H_{z,ijk-1/2}}{\Delta x} = S_{ijk}, \end{aligned} \quad (28)$$

where we have added the source term S_{ijk} (defined below).

We still need to recombine the gas energy and potential energy parts of the flux to form a single equation for the gas energy. This is

accomplished through the transformation

$$\begin{aligned} H_E &\rightarrow H_E + H_\phi \\ H_\phi &\rightarrow 0, \end{aligned} \quad (29)$$

where H_E refers to the gas energy flux and H_ϕ refers to the potential energy flux.

The source term, S , is equal to the RHS of the evolution equations 1, 2, 3, 4, and 10,

$$S = \begin{bmatrix} 0 \\ \rho \mathbf{g} + \Omega \times \mathbf{s} \\ \frac{1}{2} (\phi \frac{\partial}{\partial t} \rho - \rho \frac{\partial}{\partial t} \phi) + \Omega \times (\rho \mathbf{x} \times \mathbf{g}) \\ 0 \\ 0 \end{bmatrix} \quad (30)$$

Equations (28) and (29) account for all terms of equations (1) through (3) and equation (10) except for the $\frac{\partial}{\partial t} \frac{1}{2} \rho \phi$ term on the LHS of equation (3). This term is accounted for by updating the total energy every time the FMM solver is called to update the potential,

$$E_{\text{after}} = E_{\text{before}} + \frac{1}{2} \rho (\phi_{\text{before}} - \phi_{\text{after}}), \quad (31)$$

where the 'before' and 'after' subscripts refer to before and after the FMM solver is called.

The $\frac{d}{dt}$ operator from equation (28) is discretized using the third-order total variation diminishing Runge–Kutta integrator of Shu & Osher (1989). The time-step size, Δt , is chosen by the Courant–Friedrichs–Lewy (CFL) condition

$$\Delta t = \eta_{\text{CFL}} \max_{\text{all}} \left[\frac{\Delta x}{a_{x|y|z}} \right], \quad (32)$$

where the maximum is taken over all the computational cells in the domain and η_{CFL} is a positive dimensionless constant less than $\frac{1}{2}$. For the binary simulations presented below, $\eta_{\text{CFL}} = \frac{4}{15}$.

A schematic representation of the execution of the steps is shown in Algorithm 1.

2.4 Adaptive mesh refinement

OCTO-TIGER uses an oct-tree based AMR. Each node of the oct-tree has associated with it a single $N \times N \times N$ subgrid and is either fully refined with eight child nodes (an interior node) or not refined at all (a leaf node). For the simulations in this paper, the subgrids' interior size is $8 \times 8 \times 8$. PPM requires a three cell boundary for a total subgrid size of $14 \times 14 \times 14$. OCTO-TIGER properly nests the subgrids, meaning there can be no more than one jump in refinement levels across subgrid boundaries.

We define an 'AMR boundary' as a subgrid boundary across which the refinement level changes. Because of proper nesting, this involves only two refinement levels, which we refer to here as the 'fine' and 'coarse' level. The fine subgrid at an AMR boundary cannot get its boundary cells from a neighbouring subgrid of the same refinement level, so it must interpolate its boundary from the coarse subgrid sharing the same boundary. To describe our AMR boundary interpolation scheme, we use an indexing system, which is aligned such that $\mathbf{x}_{i-1/2j-1/2k-1/2}^C = \mathbf{x}_{2i-1,2j-1,2k-1}^F$, where the superscripts 'C' and 'F' refer to the coarse and fine levels, respectively. The coarse sub-grid has cell centres at \mathbf{x}_{ijk} while the fine sub-grid has cell centres at half integer locations, $\mathbf{x}_{2i\pm 1/2,2j\pm 1/2,2k\pm 1/2}^F$. The slopes for our interpolation scheme for a variable u (note, the u variable used here is not the inertial frame velocity used everywhere else in this

```

Set initial conditions;
Do initial output;
while do
  for  $i \leftarrow 1$  to 3 do
    if  $i == l$  then
       $U_0 \leftarrow U$ ;
      Compute  $\Delta t$ ;
    end
    Compute  $\frac{\partial}{\partial t} U$  due to advection only;
    Compute  $\mathbf{g}$ ,  $\phi$ , and  $\frac{\partial}{\partial t} \phi$  using the FMM;
    Add the source terms to  $\frac{\partial}{\partial t} U$ ;
    Update  $U$  from  $\frac{\partial}{\partial t} U$  and  $U_0$  using the 3rd order R-K
    method;
    Update the entropy tracer using the dual energy
    formalism;
    if using floors then
      Apply floor values for density and entropy;
    end
  end
end
if time to check refinement then
  refine and de-refine the AMR mesh as needed;
  redistribute the work-load;
end
if time to output then
  output to SILO file;
end
end

```

Algorithm 1: The algorithm Octo-Tiger uses for evolution

paper) are

$$\begin{aligned}
u_{x,ijk}^C &= \minmod[u_{i+1,j,k}^C, u_{ijk}^C, u_{i-1,j,k}^C] \\
u_{y,ijk}^C &= \minmod[u_{i,j,k+1}^C, u_{ijk}^C, u_{i,j,k-1}^C] \\
u_{z,ijk}^C &= \minmod[u_{i,j,k+1}^C, u_{ijk}^C, u_{i,j,k-1}^C] \\
u_{xy,ijk}^C &= \minmod[u_{i+1,j+1,k}^C, u_{ijk}^C, u_{i-1,j-1,k}^C] \\
u_{xz,ijk}^C &= \minmod[u_{i+1,j,k+1}^C, u_{ijk}^C, u_{i-1,j,k-1}^C] \\
u_{yz,ijk}^C &= \minmod[u_{i,j+1,k+1}^C, u_{ijk}^C, u_{i,j-1,k-1}^C] \\
u_{xyz,ijk}^C &= \minmod[u_{i+1,j+1,k+1}^C, u_{ijk}^C, u_{i-1,j-1,k-1}^C].
\end{aligned} \tag{33}$$

where *minmod* is zero if the signs are opposite, or the minimum absolute value is taken if the signs are the same. The interpolation scheme is then

$$\begin{aligned}
u_{2i+q,2j+r,2k+s}^F &= u_{ijk}^C + \frac{9}{64} (\text{sgn}[q]u_{x,i,j,k}^C + \text{sgn}[r]u_{y,i,j,k}^C + \text{sgn}[s]u_{z,i,j,k}^C) \\
&+ \frac{3}{64} (\text{sgn}[q]\text{sgn}[r]u_{xy,i,j,k}^C + \text{sgn}[q]\text{sgn}[s]u_{yz,i,j,k}^C + \text{sgn}[r]\text{sgn}[s]u_{xz,i,j,k}^C) \\
&+ \frac{1}{64} \text{sgn}[q]\text{sgn}[r]\text{sgn}[s]u_{xyz,i,j,k}^C,
\end{aligned} \tag{34}$$

where $q = \pm 1/2$, $r = \pm 1/2$, and $s = \pm 1/2$ and we emphasize that the variable s does not mean the inertial frame momentum density of equation (5) and is used only here. The string ‘sgn’ means the sign of q , r , and s .

To ensure conservation across the AMR boundary, coarse fluxes on AMR boundaries are taken to be the sum of the fine fluxes through

the cell face:

$$\begin{aligned}
H_{x,i+1/2,jk}^C &= \frac{1}{4} H_{x,2i+1,2j+1/2,2k+1/2}^F \\
&+ \frac{1}{4} H_{x,2i+1,2j+1/2,2k-1/2}^F \\
&+ \frac{1}{4} H_{x,2i+1,2j-1/2,2k+1/2}^F \\
&+ \frac{1}{4} H_{x,2i+1,2j-1/2,2k-1/2}^F.
\end{aligned} \tag{35}$$

The refinement criteria for a cell to be refined is either

$$\frac{\Delta x}{\rho} \frac{d\rho}{dx_i} > 0.1, \tag{36}$$

or

$$\frac{\Delta x}{\tau} \frac{d\tau}{dx_i} > 0.1, \tag{37}$$

for any i direction, where ρ is the density and τ is the entropy tracer of equation (9).

2.5 Boundary conditions

OCTO-TIGER has three boundary conditions available. With ‘inflow’ boundary conditions, the ghost cells at the edge of the physical domain are copied from the closest interior cell. ‘Outflow’ boundary conditions are the same as inflow except that the momentum is set to zero in the case where the momentum in the closest interior cell points inwards. This prevents the artificial creation of inflows. For our binary simulations, we generally use the outflow boundary condition. These boundary conditions are meant to simulate isolated systems, and do not need to be accounted for by the FMM gravity solver.

A reflecting boundary condition is available for pure hydrodynamics runs with no gravity. Although it is certainly feasible to incorporate reflecting boundary conditions into the FMM and periodic boundary conditions into both the FMM and the hydrodynamics solver, to incorporate these boundary conditions into the FMM solver is a non-trivial task. Since OCTO-TIGER does not use these types of boundary conditions to simulate binary systems, they were not implemented at this time.

2.6 Minimum values for density and entropy

Both the mass density and entropy tracers should be always positive. This condition is difficult to maintain numerically when there is a large range of densities present. For the polytropic binary simulation detailed in Section 4, we fill the regions around the stars with a density equal to 10^{-10} times the maximum density on the grid. The density can drop below this value, because the star’s gravity is constantly pulling matter from the near-vacuum region on to the stars. OCTO-TIGER also uses a third-order Runge–Kutta scheme, so even without the presence of gravity, it is not possible to know a priori the maximum time-step size needed to guarantee positivity. For these reasons, OCTO-TIGER includes options to set minimum values for the mass density and entropy tracer. These floor values are denoted ρ_f and τ_f , respectively.

Imposing a density floor results in imposing values on other variables that depend on density. We thus define the density floor scaling parameter,

$$f_\rho = 1 - \max \left[1 - \frac{\max[\rho, 0]}{\rho_f}, 0 \right] \tag{38}$$

allowing us to transform the variables according to

$$\begin{aligned} \mathbf{s} &\rightarrow \mathbf{s}f_\rho \\ E &\rightarrow Ef_\rho + \tau_f^\gamma(1 - f_\rho) \\ \tau &\rightarrow \tau f_\rho + \tau_f(1 - f_\rho) \\ \rho_m &\rightarrow \max[\rho, \rho_f] \frac{\rho_m}{\rho} \end{aligned} \quad (39)$$

This deteriorates the strict machine precision conservation of the affected variables, however, it only occurs in the near-vacuum regions.

2.7 AMR refinement criteria

OCTO-TIGER checks for refinement every $(2/\eta_{\text{CFL}})$ time-steps, where η_{CFL} is from equation (32). A window of two cells is used for refinement. This condition prevents waves from propagating through an AMR boundary before the relevant cells are checked for refinement. If a subgrid contains one or more cells flagged for refinement, the sub- is refined, converting its leaf node to an interior node with eight children. Conversely, interior nodes whose eight children are all leaf nodes are derefined if none of their cells and none of the children's cells are flagged for refinement. Subgrids are also flagged for refinement as needed to ensure the difference in refinement levels across grid boundaries is no greater than one.

For the binary simulations in this paper, we use a density based refinement criterion. If l_{max} is the maximum allowed refinement level and ρ_r is the refinement density cut-off, a cell is flagged for refinement if the maximum level l for which

$$\rho > 8^{l_{\text{max}}-l} \rho_r \quad (40)$$

holds true is greater than the cell's level of refinement. For the binary simulations ρ_r is held constant initially. Once material ejected from the binary begins to fill the grid, ρ_r is adjusted dynamically in a manner that attempts to keep the total number of subgrids at a specified level. After every refinement,

$$\rho_r \rightarrow \left(\frac{N_{\text{grids}}}{N_{\text{target}}} \right)^2 \rho_r, \quad (41)$$

where N_{grids} is the total number of subgrids and N_{target} is the desired number.

The CFL factor (equation 32) typically used by simulations is ~ 0.4 . Technically there is a limit to the maximum CFL number one should use. In three dimensions with PPM, it is $1/7$. This is because the maximum ratio between the density reconstructed at the cell's face and the cell averaged density is $7/3$. As a result, the limit should be $3/7$ in one dimension or $3/(3 \times 7)$ in three dimensions. This limit is stringent, in that it prevents any sharp transition in density from emptying a cell of its content possibly resulting in zero or negative densities. In OCTO-TIGER, negative densities are corrected by introducing a 'floor' density value in that cell (see Section 2.6), which in turn results in mass non-conservation. Pre-empting the results of Section 3.6, this can result in mass growth at the level of 0.0001 percent from one time-step to the next when highly supersonic motions are present. It is therefore important to critically appraise the value of the CFL factor to be used in each simulations, particularly if strong shocks are present. We discuss these choices and their impacts further in Section 3.6.

2.8 Code units

There are no physical constants present in the hydrodynamic equations, therefore the simulations without gravity enabled are unitless. With gravity we have a single physical constant, G . In the code, the value of this constant is set to unity. We convert between code units and physical units in the cgs system using three conversion factors, m_{cgs} , l_{cgs} , and t_{cgs} for mass, length, and time, respectively. Because the value of G is fixed at unity, we may specify two of these conversion factors, with the third being determined by the relation

$$\frac{l_{\text{cgs}}^3}{m_{\text{cgs}} t_{\text{cgs}}^2} = G_{\text{cgs}}, \quad (42)$$

where G_{cgs} is the value of G in cgs units. In OCTO-TIGER, the user specifies l_{cgs} and m_{cgs} and OCTO-TIGER calculates t_{cgs} using equation (42). When outputting the grid to file, OCTO-TIGER converts all quantities to their cgs equivalents using these unit conversion factors.

2.9 The temperature in OCTO-TIGER

OCTO-TIGER evolves the density, ρ , total gas energy density (internal and bulk kinetic energy density), E , and velocities, and derives all the other physical quantities based on these values and the EoS. The evolution equations, equations (1)–(4), do not require knowledge of the temperature. For post-processing purposes, the temperature can be computed by assigning atomic mass and atomic numbers to each mass density species and assuming a fully ionized gas according to

$$T = \frac{(\gamma - 1)\rho\varepsilon}{k_B \sum \frac{\rho_m(1+N_{Z,m})}{m_H N_{A,m}}}, \quad (43)$$

where $N_{A,m}$ and $N_{Z,m}$ are the atomic mass and atomic numbers of the m th species and the specific internal energy ε , comes from equation (11).

2.10 The C++ standard library for concurrency and parallelism (HPX)

OCTO-TIGER is parallelized for distributed systems using the C++ Standard Library for Concurrency and Parallelism (HPX; Daiß et al. 2019; Heller et al. 2019; Kaiser et al. 2020). HPX is an open source C++ Standard Library for Concurrency and Parallelism and is within the class of the so-called asynchronous many-task AMT runtime systems.

Another AMT utilized in astrophysics simulation, e.g. ChaNGa (Jetley et al. 2008) or Enzo-P (Bordner & Norman 2012), is Charm++ (Kale & Krishnan 1993). We focus on the comparison of these two AMTs in this paper; for a more comprehensive review for various AMTs, we refer to Thoman et al. (2018). The commonality of HPX and CHARM++ is the usage of the same concepts for 'parallelism' and 'concurrency'. The distinctness of HPX is that it fully conforms to the C++ 17 ISO standard (The C++ Standards Committee 2017) and implements proposed features of the upcoming C++ 20 ISO standard (The C++ Standards Committee 2020). This means that HPX's features that are available in the C++ standard can be replaced without changing the function arguments. From a programmer's perspective, HPX is more an abstraction of the C++ language while CHARM++ is more a standalone library. The requirement for OCTO-TIGER (distributed, task-based, asynchronous) are met by only few AMTs, and HPX has the highest technology readiness level according to this review (Thoman et al. 2018).

OCTO-TIGER takes advantage of four main features of HPX : (i) fine grained, (ii) task based parallelism through light weight user space threads, (iii) the use of C++ futures to encapsulate both local and remote work, and (iv) an active global address space (AGAS), whereby global objects are remotely and locally accessible (Kaiser et al. 2014; Amini & Kaiser 2019). These global objects reside in the memory on a given node but can be accessed remotely from any node.

Through HPX futures, OCTO-TIGER is able to overlap work with communication in a straightforward and efficient manner. For a given subgrid, the hydrodynamic and gravity computations are performed by an HPX thread. This thread spawns threads sending boundary data to sibling subgrids. HPX threads (Kaiser, Brodowicz & Sterling 2009) are lightweight and OCTO-TIGER may spawn hundreds or even thousands of threads per system thread. The subgrid thread also creates a set of futures encapsulating the boundary data it expects from its siblings. This allows the subgrid thread to sleep while waiting for its boundary data. When this data becomes available, HPX automatically wakes the thread, allowing computation to begin. The use of HPX futures in this manner allows OCTO-TIGER to overlap work with communication in a natural way. AGAS allows each node of the OCTO-TIGER oct-tree to be distributed across the system in a relatively simple manner, and each oct-tree node can access its children or its siblings using the same constructs regardless of whether a particular child or sibling resides locally or on a remote processor. Another benefit of HPX is that a unified application programmer interface (API) for local and remote functionality is provided. Thus, there is some simplification for the application programmer, since there is no need to deal with two different interfaces, like combining the two hybrid parallel approaches such as MPI and OPENMP. Note that HPX utilizes MPI for the communication, but provides an abstraction to the application programmer to hide the direct interaction with the MPI API.

To integrate acceleration cards, like GPUs, HPX provides two approaches: `HPX::compute` (Copik & Kaiser 2017) and `HPX::c1` (Diehl et al. 2018) to overlap GPU kernel execution with CPU work and networking. The GPU kernel execution returns a future allowing asynchronous integration of the GPU work into the overall asynchronous execution flow. `HPX::c1` provides features to integrate existing CUDA kernels and `textthpx::compute` automatically generates CUDA kernels from C++ code. OCTO-TIGER extends `textthpx::compute` to launch hand-written CUDA kernels. For more implementation details, we refer to Daiß et al. (2019b).

3 BENCHMARKING OCTO-TIGER

Below we carry out a number of simulations aimed at verifying and validating OCTO-TIGER. In doing so, we sometimes compare our benchmark test results with a similar test performed with FLASH (Fryxell et al. 2000), a well used hydrodynamic code often used for astronomical applications including binary interactions.

3.1 Benchmark design

We start testing the code by running pure hydrodynamic problems, namely, the shock tube problem (Section 3.2), and the blastwave problem (Section 3.3). The purpose is to validate and examine the performance of the hydrodynamic solver of OCTO-TIGER in isolation. This is an important task as this current version of OCTO-TIGER has an extended, more accurate (and more computationally demanding) hydrodynamic solver, which includes reconstruction of the cell

averaged values not only on the cell faces and edges but also on the cell vertices (see Section 2.3).

To test our gravity solver in isolation, we utilize the grid with the mass distribution of a uniform density sphere and let the code compute the gravitational potential without evolving it dynamically. The sphere is surrounded by a negligible amount of gas (see Section 3.4). OCTO-TIGER's gravity solver uses an opening angle parameter to modulate the accuracy of the gravity solution (see Section 2.2). We examine the accuracy of the computed potential with different values of this parameter as well with different resolutions.

The next test is a simulation of a polytropic structure (which could model certain types of stars) evolved over a number of dynamical time-scales. We check the structure stability over several dynamical time-scales, for different resolutions, values of the gravity solver opening angle and EoS (see Section 3.5).

Next, we test the polytrope in a wind tunnel to check how the star behaves as it moves through a hot, low density medium (see Section 3.6). We also test a rotating polytrope and we check for diffusion of the stellar rotation profile over time (see Section 3.7).

Finally, in Section 4, we simulate two polytropic stars orbiting one another in a detached configuration and then merging with a mass ratio of 0.5. We check stability of the two structures over a number of orbits with a number of resolutions.

We compare some of these tests with equivalent ones using the code FLASH (Fryxell et al. 2000). We also carry out scaling tests in Section 5.

3.2 Shock tube

To test the hydrodynamic solver, we run the Sod shock tube problem (Sod 1978). We use the conventional initial configuration of this problem: $\rho_l = 1$, $p_l = 1$, $v_l = 0$; $\rho_r = 0.125$, $p_r = 0.1$, $v_r = 0$, where l denotes the left-hand side of the discontinuity, and r the right-hand side. The variables ρ , p , and v are the gas mass density, pressure, and velocity, respectively. The gas is taken to be an ideal gas with an adiabatic index of $\gamma = 7/5$ (this value is historical as it pertained to molecular gas typical of air). Although the problem is one-dimensional in nature, we run it in three dimensions for testing. We simulated two configurations: in the first the discontinuity plane is $x = 0$, and in the second the discontinuity plane is $x + y = 0$. From the problem symmetry, planes parallel to the xy plane are identical. We discuss the results of the simulations in the next two subsections.

3.2.1 Shock front aligned along the x -axis

In this simulation, we set the initial discontinuity plane to be $x = 0$. As the simulation evolves in time, a shock wave propagates to the right of the box, along the positive x -axis, while a rarefaction wave propagates at the sound speed of the unperturbed denser gas to the left, along the negative x -axis. Between them the density discontinuity moves to the right. To show the convergence of the numeric solution to the analytical one, we run four simulations in which the grid is uniform and have a growing resolution of 64^3 , 128^3 , 256^3 , and 512^3 cells. We stop the simulation at time $t = 0.2$, when the shock fronts have not yet reached the grid boundary. In Fig. 1, top row, we plot the density and velocity in the x direction, respectively, at the end of the simulation, along a line perpendicular to the discontinuity and through the centre of the grid and compare it to the analytical solution.

The simulations, even with low resolution, nicely fit the rarefaction wave. Around the discontinuity, there are cells that underestimate the

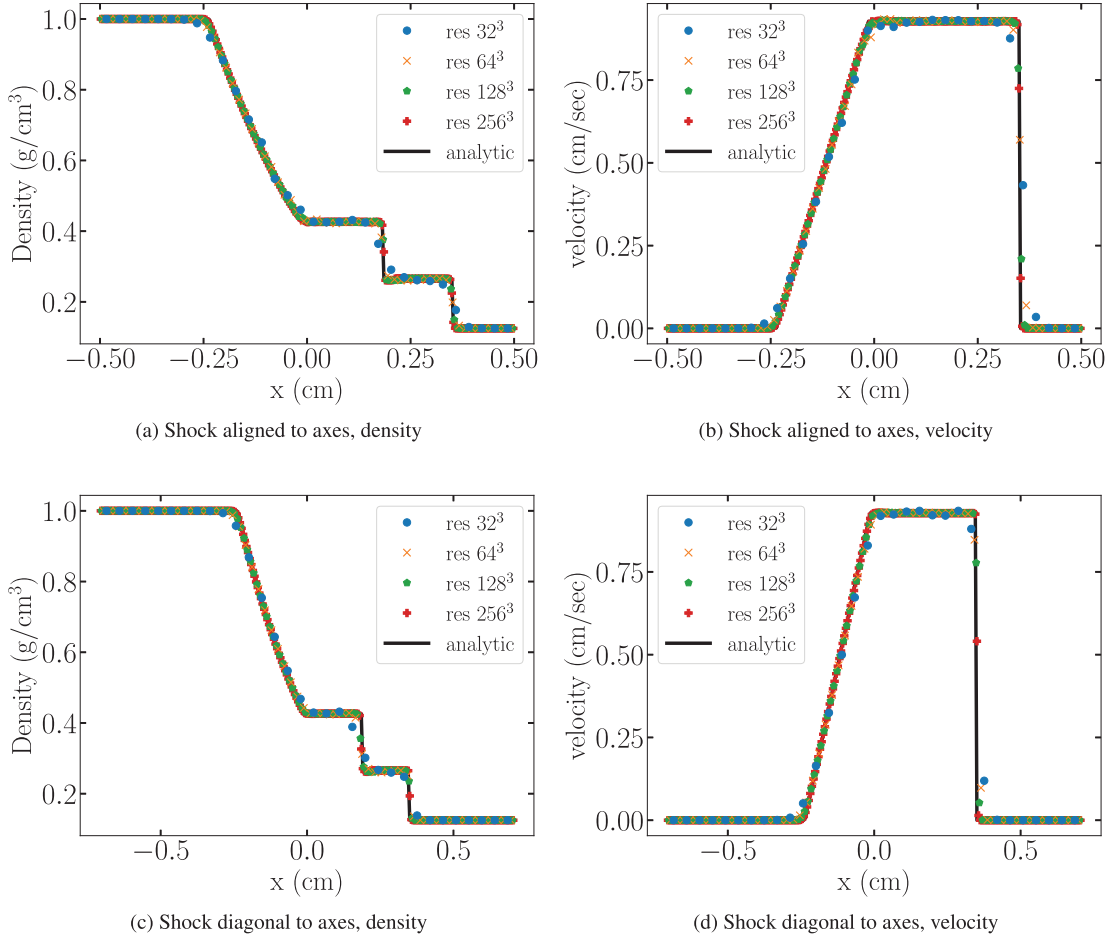


Figure 1. Comparison in density and velocity between the numerical results with higher resolution of OCTO-TIGER simulations and the exact solution for the shock tube at $t = 0.2$. Top row: The density and velocity are plotted in the x direction, respectively, at the end of the simulation, along a line perpendicular to the discontinuity and through the centre of the grid. Bottom row: The density and the parallel component of the velocity are plotted along the line that starts from the bottom left corner and ends at the upper right corner (in the xy plane).

density behind the shock and overestimate the density in front of the shock. This discrepancy diminishes with higher resolution. We ran, in addition, two simulations in which the shock is aligned to the y -axis and the z -axis. We find the same behaviour along the direction of the shock. The velocities perpendicular to the normal direction of the plane of discontinuity vanish everywhere as they should.

Overall, the Sod problem with a shock aligned to an axis shows agreement between the OCTO-TIGER simulation and the analytical solution. As expected, the numerical solution approaches the analytical one with higher resolution.

3.2.2 Shock front aligned diagonally

To check the effect of a discontinuity that is not aligned along an axis, we fixed the discontinuity plane at an angle of 45 deg from the x -axis (the $x + y = 0$ plane). A strong shock front propagates towards the higher x and y values, to the less dense upper right corner on the xy plane. A discontinuity propagates in this direction as well but with lower speed. A rarefaction wave propagates toward the bottom left corner at the sound speed. In this configuration, the wavefronts do cross the grid boundary. As the simulation evolves, an increasingly larger part of the waves encounters the grid boundary. In principle, the analytical solution is only valid in the regions where

the wavefronts did not reach the boundaries, e.g. diagonally from the bottom left corner towards the upper right corner.

In Fig. 1, bottom row, we plot the density and the parallel component of the velocity along the line that starts from the bottom left corner and ends at the upper right corner (on the xy plane) of four OCTO-TIGER simulations that differ only in their resolutions. Across this line, we notice similar behaviour to when the shock was aligned along the x -axis. The simulations reproduce accurately the rarefaction wave. Around the discontinuity, some cells underestimate the density behind the shock and overestimate the density in front of it, a discrepancy that diminishes with higher resolution.

It is, however, interesting to understand the features that appear at the opposite corners as a function of particular boundary conditions. For this test, we run two OCTO-TIGER simulations with 256^3 cell resolution that differ only in their boundary conditions. We also compare the OCTO-TIGER simulations to an identical FLASH (version 4.6) simulation. We run FLASH with the split hydro solver that uses the PPM method. The first boundary condition termed ‘outflow’ in FLASH and ‘inflow’ in OCTO-TIGER (see Section 2.5), means gas can inflow back to the grid. The second boundary condition, called ‘diode’ in FLASH and ‘outflow’ in OCTO-TIGER means no inflow is allowed.

In Fig. 2, we present slices along the xy plane ($z = 0$) of the difference in the x components of the velocity between the




Figure 2. OCTO-TIGER versus FLASH for a resolution of 256^3 and an angle of 45 deg for the shock tube. Difference in the x -velocities between simulations and analytical solution at the end of the simulation, time $t = 0.2$. The boundary condition in the top row is outflow without material inflow (diode), while in the lower row it is an outflow condition that allows material to inflow back to the simulation domain (outflow).

simulations and the analytical solution, which assumes that there is no boundary. Despite the excellent agreement with the analytical solutions (convergence is first-order near the shock front and second order at the rarefaction wave), both FLASH and OCTO-TIGER simulations show similar boundary features at the upper left and bottom right corners. At those corners, the density is smaller than the analytical solution due to gas that escapes from the grid. The gas in the upper left corner escapes easily through the $y = 0.5$ boundary and accelerates in the negative x -direction, but it also encounters the gas to its right and decelerates in the positive x -direction. The exact opposite happens at the bottom right corner. For the same reason, the differences in the y -velocities have exactly the same mirror picture. The diode boundary, by virtue of setting to zero the momentum parallel to the boundary component, does not allow velocities in the boundary direction to evolve. The rarefaction wavefront diverts near the boundary and propagates in a perpendicular direction to the boundary. This creates a square pattern at the corners.

3.3 Sedov–Taylor blast wave

The Sedov–Taylor blast wave test (Sedov 1946) has an analytical solution in three dimensions. The shock wave it produces is much stronger than that produced by the usual Sod shock tube and its spherical geometry provides a stringent test of the hydrodynamics. The initial conditions are a constant density medium with a value of 1 code unit. The internal energy density is setup so as to be described by a delta function: we approximate this in OCTO-TIGER, in a manner similar to FLASH, namely by setting the internal energy density to $E_0/(160\Delta x^3)$ for the 160 computational cells satisfying $|x| < 3.5$, and setting its value to 1×10^{-20} for all other cells. We ran this

model using AMR with 2, 3, 4, 5, 6, and 7 levels of refinement and the usual 8^3 base grid. The finest grid cell size of the model with 2 levels is 3.1×10^{-2} code units and the finest grid cell size of the model with 7 levels is 9.8×10^{-4} code units.

Density slices at $t = 0.25$ code units are shown for the models with 6 and 7 levels in Fig. 3. We also show the AMR grid structure for the run with 7 levels. In Fig. 3(d), we show the difference between the computed density and the analytical solution as a function of fine grid cell size. The blast wave is a difficult problem to obtain convergence as it requires high resolution to resolve the shock front. OCTO-TIGER obtains slightly better than first-order convergence from the run with six level of refinement to that with seven levels.

3.4 Uniform static sphere

To test the performance of the gravity solver that solves the Poisson equation, we initiate a problem with a static uniform density sphere. This problem has a simple analytical solution and is widely used to test solvers for self-gravitating fluids (Mott et al. 2002).

In this test, we place a sphere with a radius of 0.25 code units at the centre of a cubic grid with a length of 1 code units. The total mass inside this sphere is 1 code unit. We fill the grid outside the sphere with a low density of 1×10^{-10} code units. We used a uniform grid with four resolutions: 64^3 , 128^3 , 256^3 , and 512^3 . To demonstrate the effect of changing θ , the opening-angle parameter (see Section 2.2), we carried out two additional simulations with resolution of 128^3 and 256^3 , where $\theta = 0.35$ instead of the default value of 0.5. The OCTO-TIGER simulations reproduce very accurately the analytical gravitational potential. The residuals of the potential

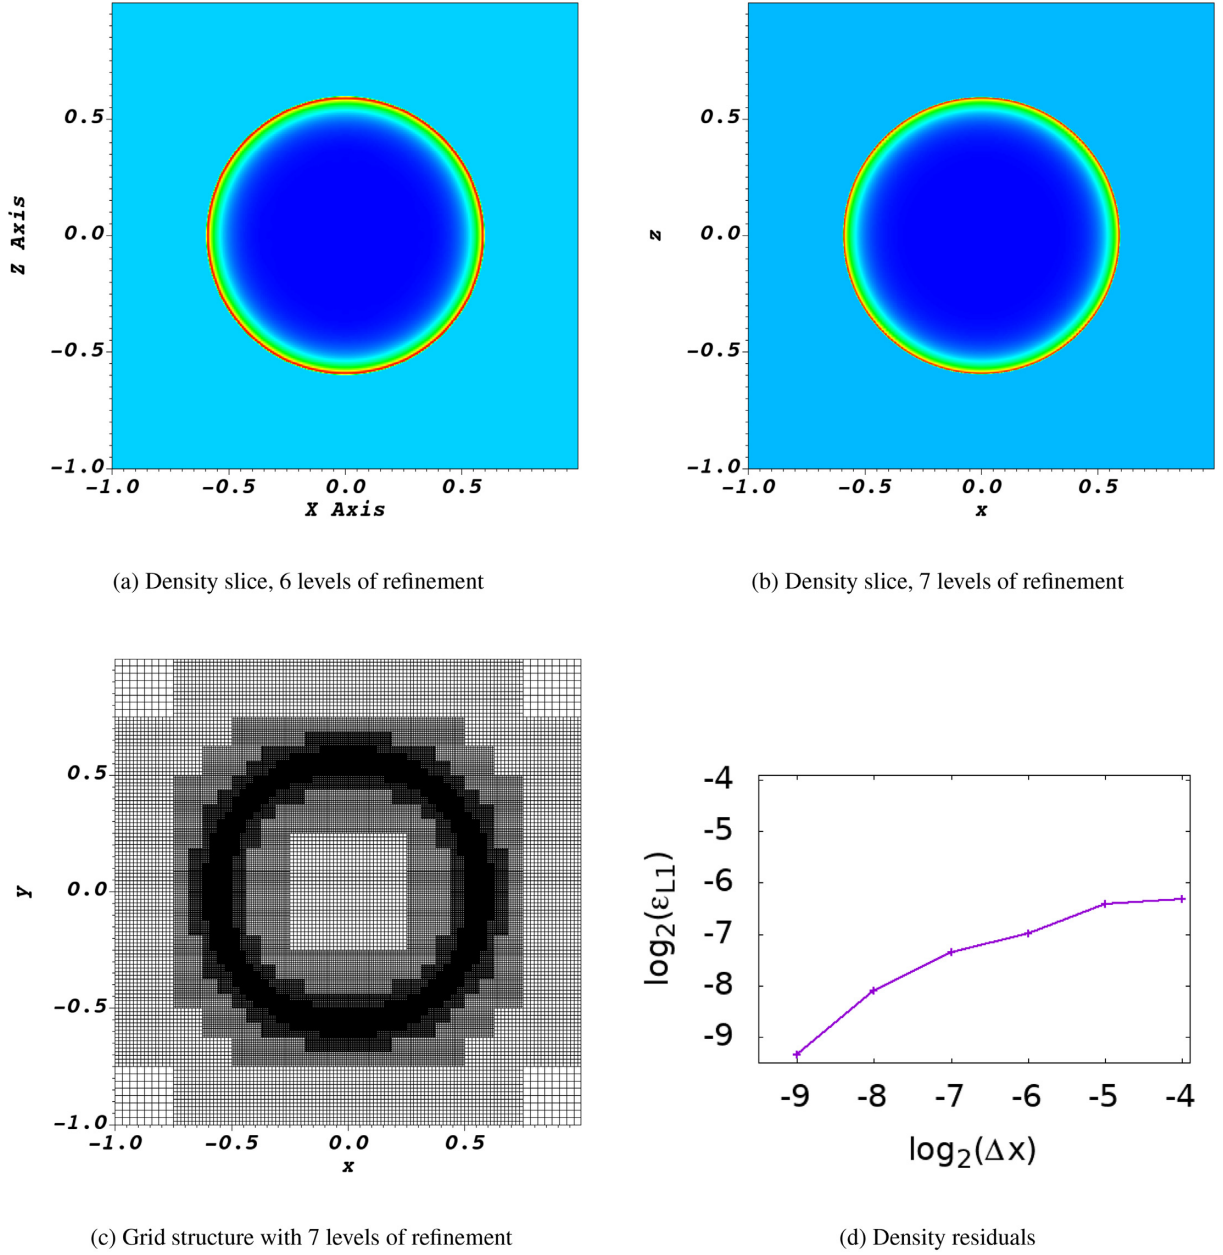


Figure 3. The Sedov–Taylor blast wave test. Density slices using OCTO-TIGER with 6 (a) and 7 (b) levels of refinement at $t = 0.25$ code units. Panel (c): the AMR grid structure of the simulation with 7 levels. Panel (d): the absolute deviation of the density from the analytical solution as a function of grid cell size for six simulations carried out with a maximum of seven levels of refinement to a minimum of 2.

on the equatorial plane

$$\epsilon = \left| \frac{\phi - \phi_{\text{analytic}}}{\phi_{\text{analytic}}} \right|. \quad (44)$$

are plotted in Fig. 4 for four of the OCTO-TIGER simulations. A square pattern is apparent for the residuals in the OCTO-TIGER simulations with the maximum residuals appearing at the corners of a square that encloses the sphere. These larger residuals do not decrease with higher resolution, something that is expected of the FMM method. The FMM computes the multipole expansion of the potential between grid cells at the same refinement level as each other. Adding extra refinement to those levels only increases the quality of the solution between cells at the extra refinement level. The cells at the coarser levels still interact with each other in the same way, and the same

expansions are passed to the more refined levels. The only way to cause cells to interact with each other at finer levels of refinement is to decrease the opening criterion.

We tested whether a smoother transition of the sphere to the background density can lower the residuals by carrying out simulations of two other density distributions (which have known analytical solutions), one where the density of this sphere decreases to the ambient value by a parabolic dependence with radius (continuous but not differentiable), and a second where the density distribution obeys an $n = 1$ polytropic profile (continuous and differentiable). We find that regardless of the smoothness of the density decrease, high residual values remain at the corners. To decrease the high residuals at the corners, a lower value of θ needs to be set. We also tried to simulate an off-centre sphere and a range of additional sphere sizes.




Figure 4. Uniform static sphere test. Maps of the residuals of the gravitational potential, ϵ , at the equatorial plane $z = 0$ as a function of resolution, modelling code, and opening angle θ .

Changing the sphere location does not affect the residuals. However, increasing the sphere radius does result in lower maximum residual values. We conclude that the maximum values are related to how much mass is concentrated in a given volume.

For comparison, in Fig. 4, we present equatorial maps of several 128^3 and 256^3 cell FLASH simulations that used two gravity solvers: multigrid (panel e) and BHTree (panel f). Here, we see that each solver has a different residual shape.

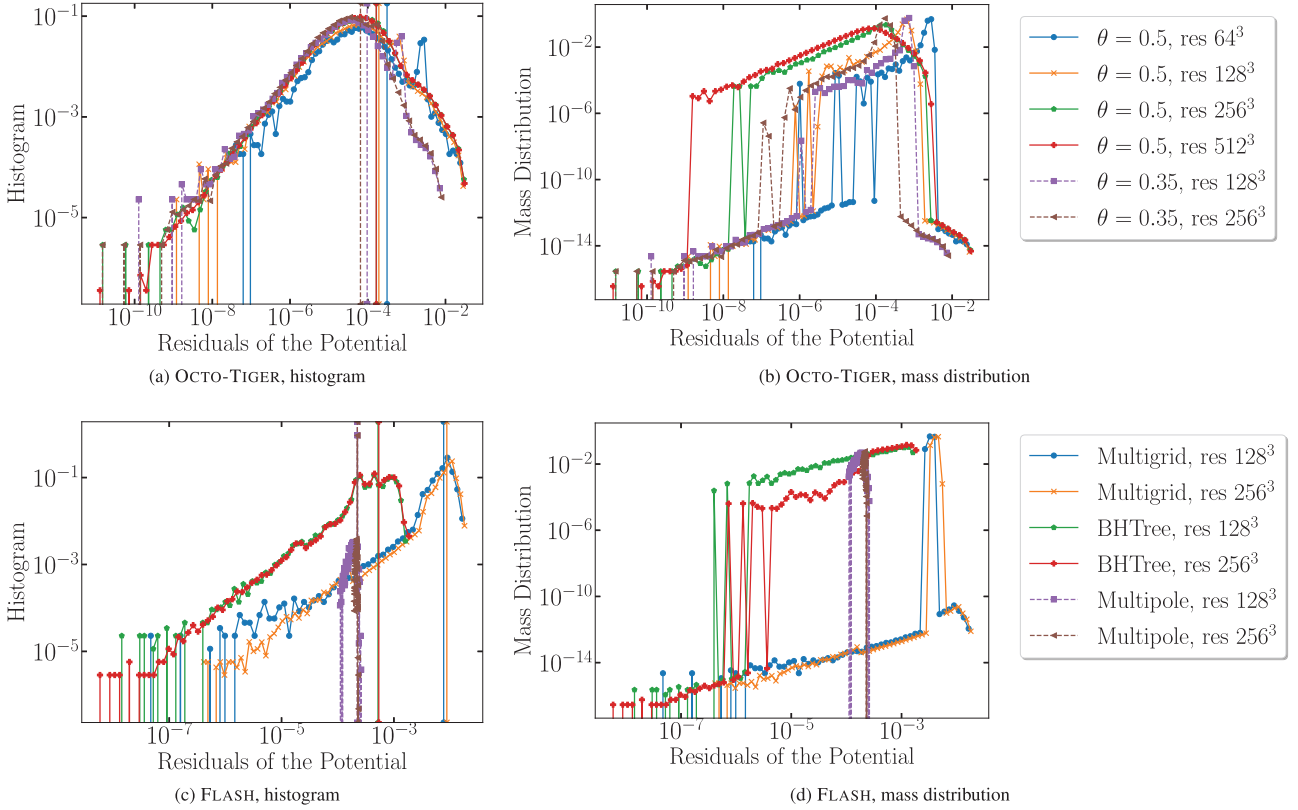


Figure 5. Uniform static sphere test. Histograms and mass distributions of the potential’s residuals of OCTO-TIGER and FLASH simulations. The vertical lines in panels (a) and (c) mark the residuals’ mean value for each simulation.

In Fig. 5, we plot histograms of the residuals by volume (panels a and c) and by mass (panels b and d). Fig. 5(a) shows that the residuals’ maxima do not decrease by increasing the resolution; however, the residuals’ mean and minimum values do. By increasing the resolution, the distribution is stretched to lower residual values. The notable second peak of the histogram for the lower resolution simulations practically disappears for the two higher resolutions. The most common residual value is consistently around 10^{-4} irrespective of resolution.

In Figs 5(c) and (d), we present the residuals’ histograms and mass distributions of the FLASH simulations, where we used three commonly used FLASH gravity solvers: Multigrid, BHTree, and Multipole. The FLASH multipole gravity solver has the lowest maximum residuals. The 128^3 simulation that used the multipole solver has a lower mean value $\simeq 2.23 \times 10^{-4}$, than every one of the other FLASH simulations (including the high-resolution simulations). Only the OCTO-TIGER simulations have lower mean residuals: 1.93×10^{-4} , 1.67×10^{-4} , and 1.65×10^{-4} for the 128^3 , 256^3 , and 512^3 simulations, respectively. The FLASH multipole solver easily produces an accurate solution to this problem because it has a spherical symmetry. Including additional terms, besides the monopole, does not contribute to the numeric solution in the FLASH multipole solver. The maximum and mean values of the residuals slightly increase in the high-resolution FLASH simulations. The OCTO-TIGER simulations show overall low mean residuals and low most common residuals (the peak in the distributions both with respect to volume and mass).

We summarize the residuals’ minimum, maximum, and mean values for the OCTO-TIGER and FLASH simulations in Tables 1 and 2, respectively. Table 1 shows that in OCTO-TIGER, a higher resolution decreases the mean residual value, up to a certain resolution, where

the mean residual value converges to some low value. To decrease the mean residual further, one must set a lower θ value.

From the mass distributions (Fig. 5b), we see that the cells that have the highest residual values actually contain a negligible amount of mass. Namely, the cube corners that have high residuals reside outside the uniform density sphere, where the density is very low. In addition, increasing the resolution results in having most of the mass in lower residuals. The residual value, that has the maximum mass, shifts from 3×10^{-3} in the 64^3 simulation to 10^{-4} in the 512^3 simulation, and the mass distribution becomes flatter towards smaller residuals with higher resolution. We infer that these regions with high values of the potential residuals will have minimum effect on the evolution of a simulation, once we evolve it, because of their low total mass.

As can be seen from Figs 4 and 5, by decreasing the opening-angle parameter, θ , we can increase the level of accuracy in the OCTO-TIGER simulations even more. The residuals altogether are reduced, including the maximum residual values at the corners. With such low value of θ , we get the lowest mean residual values among all the simulations even at the lower resolution of the 128^3 cells. As with the simulations with $\theta = 0.5$, improving the resolution does not reduce the residual maximum, but does improve the minimum and mean values. By increasing the resolution, the main peak is shifted to lower residual values of around 4×10^{-5} . The mean residual value in the high-resolution simulation indicates a deviation of 0.0065 per cent from the analytical value. We note, though, that the peak in the mass distribution of the residuals remains approximately the same for the $\theta = 0.35$ simulations and for the $\theta = 0.5$ simulations (compare the peaks of the purple square lines and the yellow cross lines or the peaks of the green circle lines with the brown triangle

Table 1. Gravitational potential residuals of the static uniform sphere in OCTO-TIGER. OCTO-TIGER uses the fast multipole method to solve for the gravity, which has an opening angle parameter θ . Full distributions appear in Fig. 5.

N_{cells}	64^3	128^3	256^3	512^3	128^3	256^3
θ	0.5	0.5	0.5	0.5	0.35	0.35
Max	2.66×10^{-2}	3.23×10^{-2}	3.49×10^{-2}	3.61×10^{-2}	0.82×10^{-2}	0.90×10^{-2}
Mean	3.13×10^{-4}	1.93×10^{-4}	1.67×10^{-4}	1.65×10^{-4}	9.6×10^{-5}	6.5×10^{-5}
Min	3×10^{-8}	1×10^{-9}	1×10^{-11}	1×10^{-11}	1×10^{-10}	1×10^{-11}

Table 2. Gravitational potential residuals of the static uniform sphere in FLASH. Full distributions appear in Fig. 5.

Solver	BHTree	BHTree	Multigrid	Multigrid	Multipole	Multipole
N_{cells}	128^3	256^3	128^3	256^3	128^3	256^3
Max	0.18×10^{-2}	0.20×10^{-2}	1.81×10^{-2}	1.96×10^{-2}	2.6×10^{-4}	2.3×10^{-4}
Mean	5.21×10^{-4}	5.38×10^{-4}	7.691×10^{-3}	8.759×10^{-3}	2.23×10^{-4}	2.24×10^{-4}
Min	1×10^{-8}	5×10^{-9}	4×10^{-8}	4×10^{-7}	1.1×10^{-4}	1.9×10^{-4}

lines in Fig. 5b). Reducing θ has a higher computational cost (see Section 5). This compromise between the accuracy of the simulation and the running time should be taken into account when simulating a problem.

3.5 Stationary star

In this test, we set up a polytrope with an index of $n = 3/2$. The stellar diameter is the same as we have used in the uniform static sphere test, equal to half of the domain size, and the polytrope’s centre coincides with the centre of the domain. A polytrope can be scaled to model different types of stars. We show two such scalings in Table 3, one for a low-mass, fully convective, main-sequence star and a second one for a low-mass WD. Outside the star, we fill the domain with gas with a density that is 10 orders of magnitudes smaller than the star’s central density. While the star itself is in good hydrostatic equilibrium, the outside medium is not in equilibrium and it starts free falling on to the surface of the star as the simulation starts. The falling of the gas together with the diode boundary condition, which prevents inflow, create outer regions with very low densities. To avoid the creation of a vacuum, whenever cell densities become lower than the threshold floor density of 10^{-15} times the central density, we set the density to be that value. In addition, to reduce shocks due to supersonic in-falling gas we give the ambient medium a high internal energy (and hence temperature). We list the ambient medium’s properties in Table 3.

Any small perturbation from hydrostatic equilibrium will induce an oscillation with polytrope eigenfrequencies and modes. Hurley, Roberts & Wright (1966) computed numerically the fundamental modes, as well as the first and second harmonic modes of several polytropes with different adiabatic indices. They found that a pulsating $n = 3/2$ polytrope has a fundamental frequency of

$$\omega_F^2(n = 3/2; \gamma_{\text{ad}} = 5/3) = 0.3764 \frac{8\pi G \rho_c}{5}. \quad (45)$$

The pulsation periods are of the order of the free-falling time of the star. For the low-mass main-sequence star model, for example, the fundamental period, T_F , is 24 min, while for the WD model the fundamental period is 18 s.

We compare the oscillations observed in the simulations with the analytical solutions. We also monitor the diffusion of the outer layers of the star and the behaviour of the low-density medium that

surrounds the structure and we verify the conservation of some basic physical quantities.

We ran six simulations in total, two OCTO-TIGER uniform grid simulations with resolutions of 128^3 and 256^3 cells and one OCTO-TIGER AMR simulation with base resolution of 128^3 and one level of refinement that increases linearly the resolution by a factor of 2. We refine based on a criterion of density (see Section 2.4). These three simulations use the OCTO-TIGER FMM gravity solver, with an opening angle parameter $\theta = 0.5$. We also ran a fourth uniform grid OCTO-TIGER simulation with resolution 128^3 , and $\theta = 0.35$.

These four simulations used an ideal gas EoS. The fifth simulation, a uniform grid of 128^3 resolution and $\theta = 0.5$, used a polytropic EoS (see equation 14). The sixth simulation is carried-out with FLASH on a uniform grid resolution of 128^3 , using the BH-tree gravity solver.

In OCTO-TIGER, we solve for each cell in the domain of the Lane–Emden equation to obtain the polytrope, while in FLASH we interpolated a one-dimensional polytropic solution into the three-dimensional grid. This results in small differences between the initial states of the OCTO-TIGER and FLASH simulations. In Fig. 6(a), we show the AMR grid in a temperature slice at the equatorial plane of the AMR OCTO-TIGER simulation. Each cube represents a subgrid containing 8^3 equal volume cubic cells. As we refined by the density criterion, only the inner region of the star is refined to the maximum level. In Fig. 6(b), we plot the initial density profile of the stars in both OCTO-TIGER and FLASH.

In Fig. 7(a), we plot the central density of the star, divided by the initial central density, over time (in units of T_F , the fundamental pulsation period of our polytrope). As expected, in all simulations, the central density of the star shows clear oscillations. In the OCTO-TIGER simulations, the central density oscillates around some converging value, while in FLASH, there is a slow decrease with time overlaid on the central density oscillations. By setting a lower value for the opening angle parameter of the gravity solver ($\theta = 0.35$, red crosses in Fig. 7), we can lower the amplitude for these oscillations.

Although the oscillation amplitude might affect the amount of noise in the solution, the clear-cut test is whether the oscillation frequencies are aligned with the fundamental frequency predicted by theory. The assumption is that both the initial state and the numerical solving scheme inaccuracies will induce random noise that will oscillate at the fundamental frequency of the polytrope. For that, we plot the Fourier transform of the stars’ central density (Fig. 7b), normalized by the initial pulsation amplitude, in units

Table 3. Examples of how our polytropic model of Section 3.5–3.7 can be scaled to represent stars of different types. We assume $\mu_{\text{He}+2} = 4/3$, and $\mu_{\text{CO}} = 2$, $\mu_{\text{H}} = 1$, $\mu_{\text{H}+} = 1/2$, and for calculating the temperature for the main-sequence star centre, WD centre, main-sequence star surface, and WD surface, respectively. To calculate the medium’s temperature, we assume $\mu_{\text{H}+} = 1/2$. Note that T_f is the fundamental period, while T_s , T_c , and T_{medium} are temperatures.

Model	M (M_{\odot})	R	ρ_c (g cm^{-3})	p_c (dyne cm^{-2})	T_c (K)	$c_{s,c}$ (km s^{-1})	T_s (K)	$c_{s,s}$ (km s^{-1})	T_{medium} (K)	$c_{s,\text{medium}}$ (km s^{-1})	T_f (s/min)	K ($\text{erg cm}^2/\text{g}^{5/3}$)
MS	0.27	$0.25R_{\odot}$	150	1.7×10^{17}	1.8×10^7	430	6300	9.3	1.9×10^8	2300	24 min	4×10^{13}
WD	0.35	10^9 cm	10^6	2.5×10^{22}	6.1×10^8	2000	7.3×10^4	44	4.3×10^9	11 000	18 s	2.5×10^{12}

Note. Legend: MS = main sequence; WD = white dwarf; M = mass; R = radius; ρ_c = central density; p_c = central pressure; T_c = central temperature; $c_{s,c}$ = central sound speed; T_s = surface temperature; $c_{s,s}$ = surface sound speed; T_{medium} = medium temperature; $c_{s,\text{medium}}$ = medium sound speed; T_f = fundamental period; K = polytropic constant.

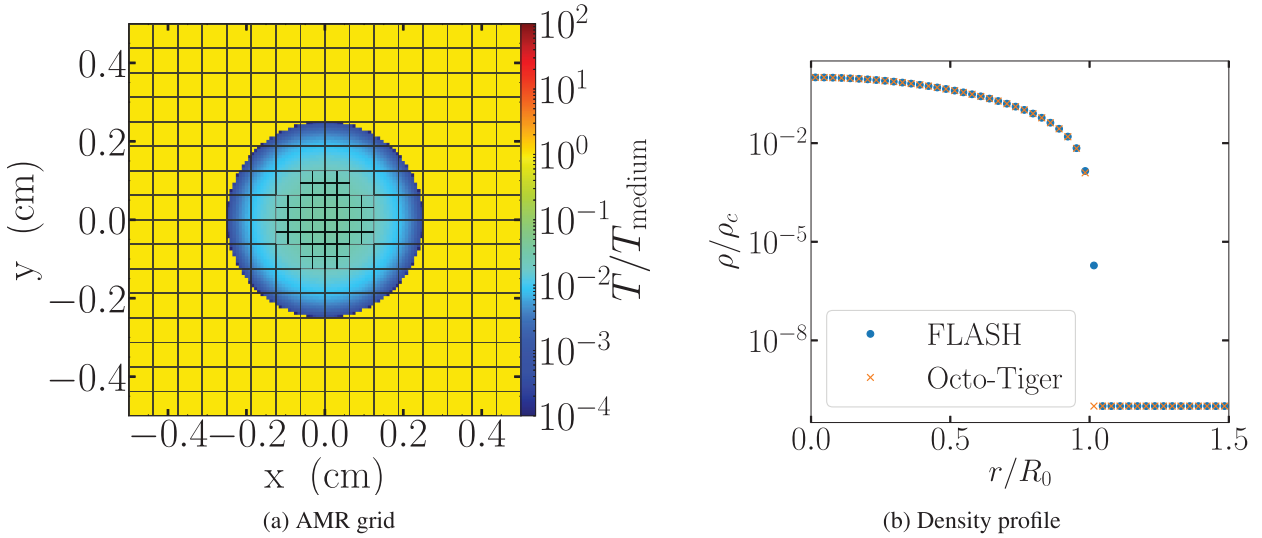


Figure 6. Stationary star initial state. (a) AMR grid in a temperature slice at the equatorial plane of the AMR OCTO-TIGER simulation. Each cube represents a subgrid containing 8^3 equal volume cubic cells; (b) the initial density profile of the stars in OCTO-TIGER and FLASH. Small differences are present due to different interpolation schemes.

of the fundamental frequency, f_f . The vertical lines in Fig. 7(b) show the fundamental frequency and the first and second harmonic frequencies, $2.15f_f$ and $3.13f_f$, respectively. The Fourier transform of the densities of all the simulations notably peak close to the fundamental frequency. Again, the star, that is evolved with a $\theta = 0.35$, reproduced most accurately the fundamental mode, with a frequency peak that deviates by only 0.5 per cent from the theoretical fundamental frequency. The other OCTO-TIGER simulations deviate by 1.5 per cent, while FLASH deviates by 3.8 per cent. The star that is evolved with a polytropic EoS, naturally suffers least from low-frequency noise.

We next plot (Fig. 7c) the deviation from conservation of mass, $\Delta M/M_0$, in the OCTO-TIGER simulations. $\Delta M = M(t) + M_{\text{out}} - M_0$, where $M(t) = \int_V \rho dV$, M_{out} , and M_0 are the mass inside the simulation domain, the mass that outflowed from the grid, and the initial mass, respectively. OCTO-TIGER calculates automatically, for every time-step, the outflow of quantities such as mass, energy, and entropy, which simplify following of every small change of those quantities. Therefore, we plot only the conservation of mass in OCTO-TIGER simulations. We see conservation at the 10^{-13} level up to a time when the density flooring adds mass to the grid. Even then the conservation, by the end of the simulation is still at the level of a few $\times 10^{-12}$.

In these simulations, we set a density floor that causes a deviation from conservation at a machine precision level. The floor density is 5 orders of magnitude less than the initial density of the ambient medium, and usually will not affect mass conservation. However,

at machine precision level, a deviation that grows approximately linearly with time appears. This can be explained by considering the following. As a consequence of ambient medium gas falling on to the star surface, the density outside the star decreases with time. Eventually, the density will decrease below the density floor threshold at the outer regions, which is then filled with density at the floor value. This happens after approximately 20, 25, and 35 fundamental periods in the low-resolution atmosphere with ideal gas EoS simulations, polytropic EoS simulation, and the high-resolution simulation, respectively. If the region that is being filled with a floor density is a fraction α of the domain volume, we will get a deviation from mass conservation of roughly $\alpha V \rho_{\text{floor}}/M_0 = \alpha \rho_{\text{floor}}/\bar{\rho}_0$, where $\bar{\rho}_0$ is the initial mean density in the simulation’s domain. In our simulations $\rho_{\text{floor}}/\bar{\rho}_0 \simeq 9 \times 10^{-14}$. If only one layer of cells at the grid boundary is being filled with the floor density then $\alpha \simeq 0.05$. If, in addition, the cells will be refilled every five time-steps, after $\sim 20\,000$ time-steps a deviation of 2.2×10^{-11} from conservation of mass will occur. In the simulations themselves, fewer cells are refilled less frequently so the deviations are smaller. Additionally, as shock heating is eliminated in the polytropic EoS simulation, the ambient gas continuously falls on to the star’s surface without disturbance, which results in a bigger volume of flooring, and in a somewhat higher deviation from conservation of mass (purple line).

In Fig. 7(d), we show the total mass enclosed in a sphere with a radius of the initial star’s radius. Here, we see that FLASH loses 1 per cent of the stellar mass in 40 pulsations, while the OCTO-

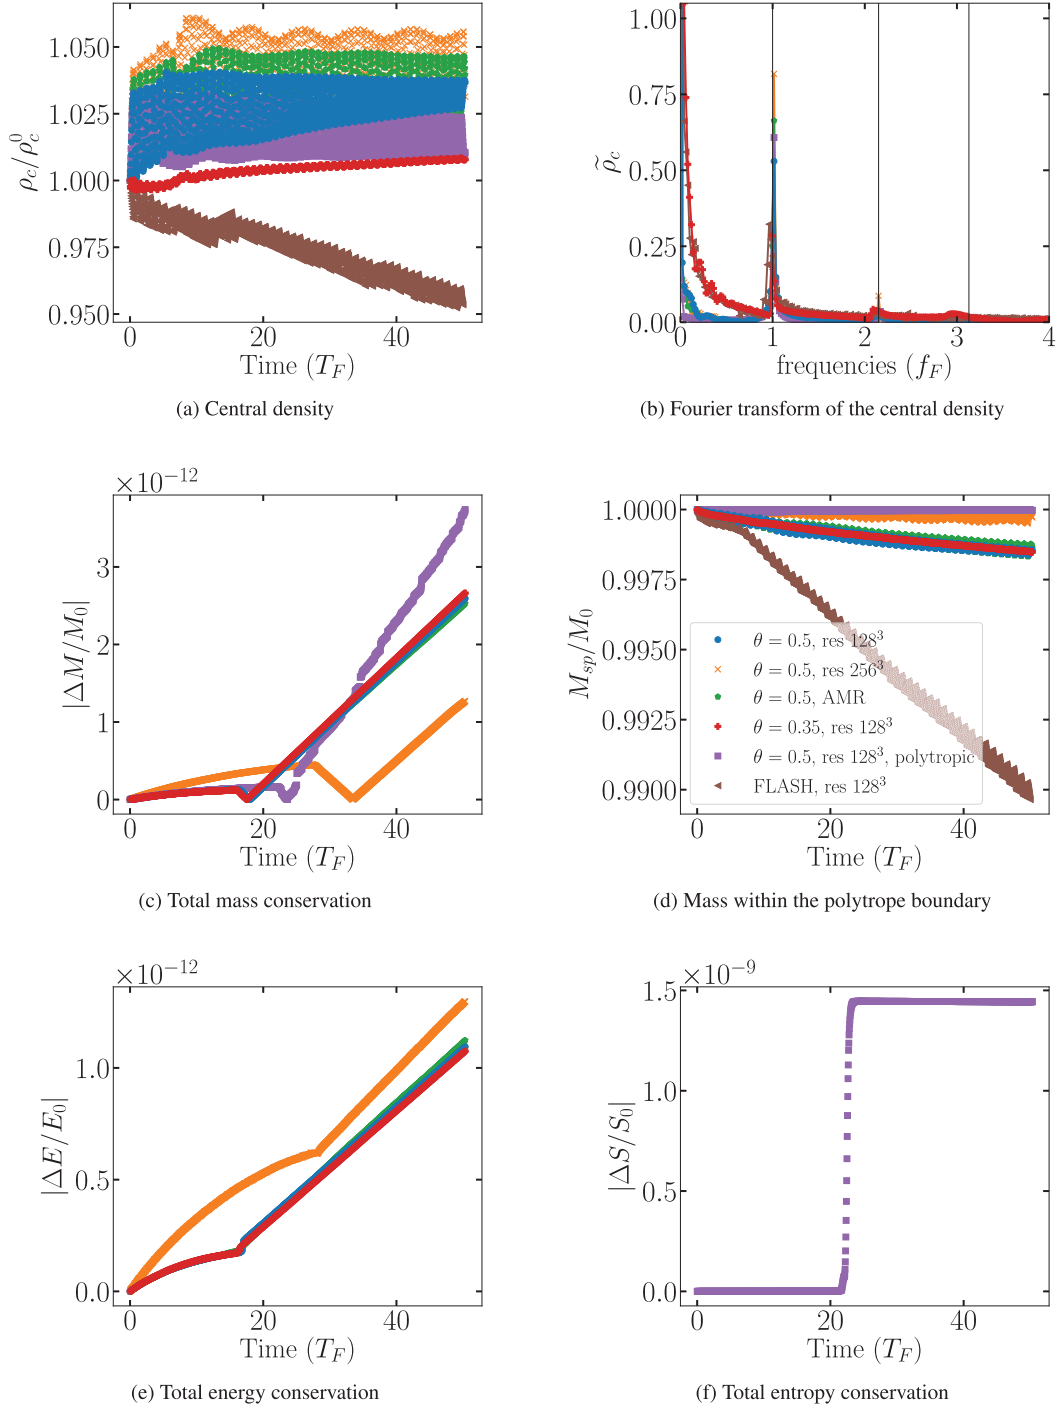


Figure 7. The stationary polytrope benchmark test. (a) Central density over time divided by the initial central density. (b) The Fourier transform of the central density pulsations normalized by the initial pulsation amplitude. The vertical lines show the fundamental frequency and the first and second harmonic frequencies. (c) Conservation of mass (only for OCTO-TIGER simulations). (d) Total mass enclosed in a sphere with a radius of the initial polytropic radius. (e) Conservation of energy (only for OCTO-TIGER simulations with ideal gas EoS). (f) Conservation of entropy (only for the polytropic EoS OCTO-TIGER simulation). T_F is the fundamental period of a polytrope with $n = 3/2$ and $\gamma = 5/3$.

TIGER simulations lose at most 0.2 per cent of the mass flowing out. The lower resolution, polytropic EoS, and the higher resolution simulations perform best.

In Fig. 7(e), we show the deviation from conservation of energy $\Delta E/E_0$, where $\Delta E = E(t) + E_{\text{out}} - E_0$, and $E(t) = \int_V (E + \frac{1}{2}\rho\phi) dV$, E_{out} and E_0 are the energy inside the simulation

domain, the energy that outflows from the grid and the initial energy, respectively. In Fig. 7(f), we show the deviation from conservation of entropy $\Delta S/S_0$, where $\Delta S = S(t) + S_{\text{out}} - S_0$, and $S(t) = \int_V \tau dV$, S_{out} , and S_0 are the entropy inside the simulation domain, the entropy that outflows from the grid and the initial entropy, respectively. Beside the small effect of the flooring, the ideal gas EoS simulations conserve

Table 4. Translating polytrope benchmark test. Scaled quantities: motion velocities, total simulation times, and displacement. For other scaled quantities, refer to Table 3.

Model	v_1 (km s ⁻¹)	v_2 (km s ⁻¹)	v_3 (km s ⁻¹)	t_{sim1}	t_{sim2} (s/min/h)	t_{sim3}	R_{travel} (R_{\odot})
MS	1.20	120	12 000	44 h	27 min	16 s	0.28
WD	5.6	560	56 000	33 min	20 s	0.2 s	0.016

energy at machine precision, while the polytropic EoS simulation conserves entropy at a machine precision.

Click on the link¹ to view a movie of the OCTO-TIGER, 256³, $\theta = 0.5$ simulation.

3.6 Star moving linearly in the grid

We next simulate a star as in Section 3.5, except we initialize the star with a non-zero bulk velocity, allowing it to translate at a given constant velocity through the grid. This exercise tests the degree to which the surface layers of our structure shear away due to interaction with the low-density medium permeating the background. It also tests how well a spherical star moves through our cubical grid and how well our code conserves the total, non-zero linear momentum.

We test three velocities (a) a subsonic velocity, $v_1 = 5.2 \times 10^{-4} c_{\text{s,medium}}$, where $c_{\text{s,medium}}$ is the speed of sound of the ambient hot medium (see Table 3); (b) an intermediate velocity, $v_2 = 5.2 \times 10^{-2} c_{\text{s,medium}}$; and (c) a high, supersonic velocity, $v_3 = 5.2 c_{\text{s,medium}}$. We list these three velocities scaled to a main sequence and WD models, as well as the simulation time (shorter for the faster stars) and the distance the star travels in Table 4. The distance the star travels equals to 1.1 times the star's initial radius for all simulations. We placed the star initially at $(-0.125, -0.125, 0)$ and gave it a velocity in the direction towards grid position $(1, 1, 0)$.

Each velocity regime contains a set of six simulations (as we have done for the stationary star): two uniform grid, 128³ and 256³ cell simulations, one AMR simulation (128³ cells and one level of refinement), with ideal gas EoS and $\theta = 0.5$; a uniform grid 128³ simulation with ideal gas EoS and $\theta = 0.35$; a uniform grid 128³ simulation with polytropic EoS and $\theta = 0.5$; and a FLASH uniform grid 128³ simulation. In Sections 3.6.1–3.6.3, we describe the results of the simulations.

3.6.1 Star translating at low velocity

In this regime, the star translates at a very low velocity, equivalent to a Mach number of 5.2×10^{-4} , not only with respect to the ambient medium speed of sound, but also with respect to the speed of sound of the coldest regions at the star's surface. Such low bulk velocities are a challenge for hydrodynamic codes. Low level subsonic noise that develops can advect the stellar momentum to the surrounding low-density gas, diminishing the star's initial low velocity. In addition, the star's low velocity allows us to run this test for over an hundred fundamental pulsations periods, the longest time we have run a polytrope simulation. During this time, the star should keep pulsating at its Eigenmodes as it translates through the grid.

The star in the OCTO-TIGER simulations slowly moves from the bottom left corner to the upper right corner (in the xy plane). As in the stationary star simulations, the star remains in hydrostatic equilibrium, while the ambient medium is free falling on to the

stellar surface. This slowly cools the ambient medium and creates shocks on the star's surface that heat the gas for ideal gas simulations. We performed the same analysis as we did for the stationary star. We find that the star's movement does not affect the star's pulsations. Almost identically to the stationary star in Fig. 7(a), the central density oscillates at the fundamental frequency in all OCTO-TIGER simulations, while the mean central density of the star simulated with FLASH decreases with time. Additionally, the star's mass (M_{sp}) remarkably decreases at the same rate as for the stationary star in Fig. 7(d), decreasing by 2 per cent in the FLASH simulation and by less than 0.4 per cent with OCTO-TIGER, after 100 pulsations.

In Fig. 8(a), we plot the centre of mass position as a function of time, while in panel (b) we show the deviation of the centre of mass velocity from its initial value. The star in the OCTO-TIGER simulations moves through the grid at a constant velocity, while in the FLASH simulations, the star slows down, starts moving in the opposite direction and then oscillates around the star's initial position. The OCTO-TIGER simulations deviate from the initial velocity by less than 1 per cent throughout the entire evolution.

We exploited OCTO-TIGER's capability to track the provenance of gas, to calculate the diffusion of linear momentum from the star to its environment. In Fig. 8(c), we plot the x -momentum of the diffuse medium divided by the initial *total* x -momentum in the OCTO-TIGER simulation. The initial oscillation is due to the sharp density gradient at the star's surface. Soon after the value settles at values close to zero. We mention though that the mass ratio between the ambient medium gas and the star is of the order of 10^{-8} . No comparison with FLASH is possible due to FLASH not tagging gas provenance.

We find, similarly to the stationary star test, that mass, energy, and entropy are conserved in the OCTO-TIGER simulations to excellent precision, except for small deviations due to adding mass when a cell's density drops below the minimum floor value (see Section 2.6). We additionally follow the conservation of x -linear momentum (Fig. 8d), finding a machine precision level until flooring starts operating at which point we observe a slow linear growth. The deviation, though, is minimal, less than 10^{-10} after 100 pulsation periods for all simulations except the polytropic EoS. The polytropic simulations suffer more from flooring (because shock heating is absent) and the deviation settles on a value of 8×10^{-9} .

Click on the link² to view a movie of the low velocity translating star.

3.6.2 Star translating at intermediate velocity

With a speed of Mach 5.2×10^{-2} with respect to the outside medium, the star translates at a subsonic velocity, but at a supersonic speed with respect to the star's surface. The simulation time is comparable to the star's fundamental period. This allows the star to pulsate only once while it translates at a higher speed than the previous test.

¹<https://youtu.be/4-ra6fY982Q>

²<https://youtu.be/8ArZqP9F93Y>

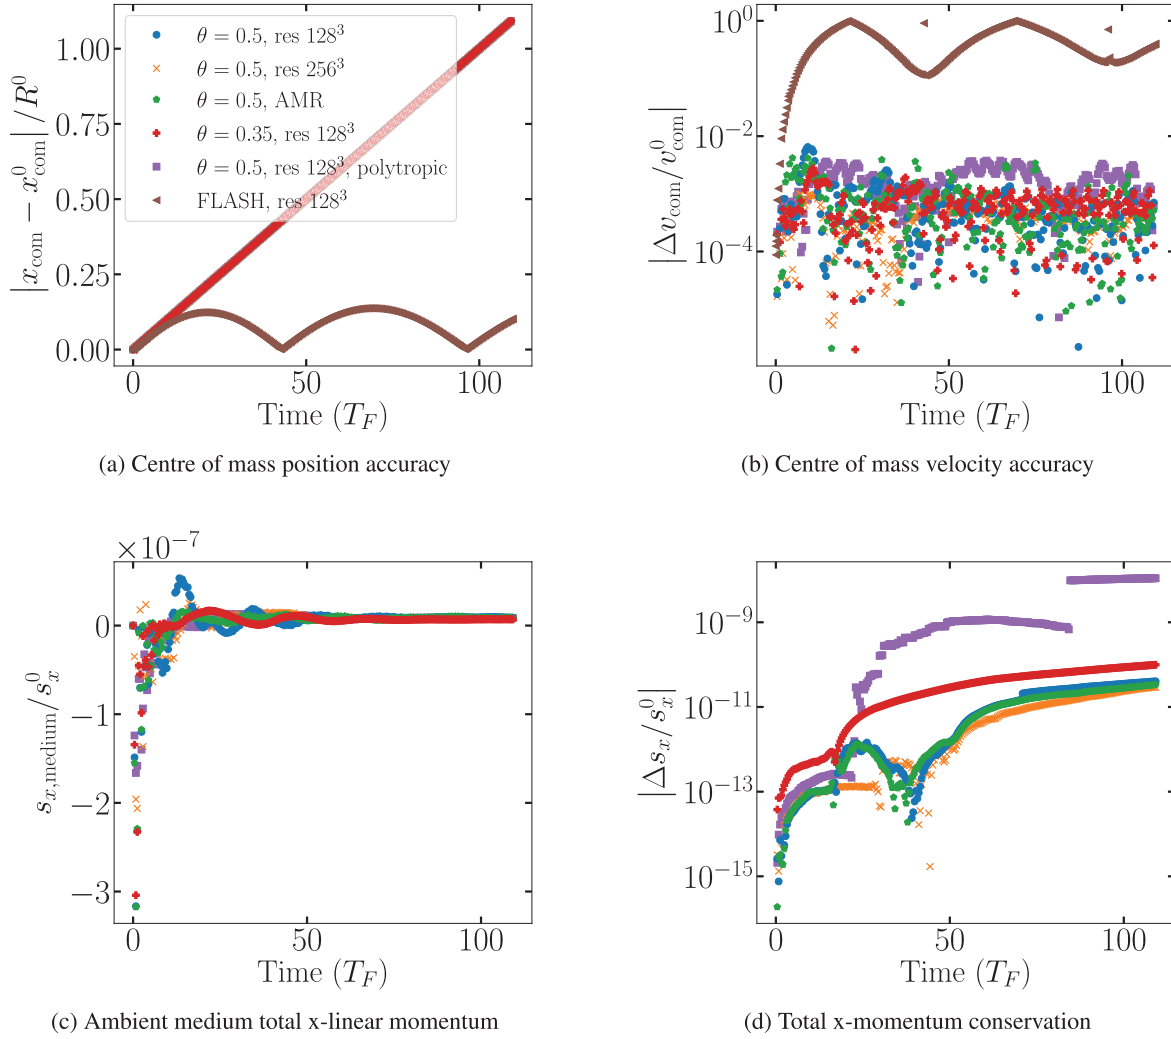


Figure 8. The translating star benchmark test with the star moving at Mach 5.2×10^{-4} (with respect to the diffuse medium). We plot various physical quantities of interest over time. Time is in units of T_F , the fundamental pulsation period.

The star’s gas diffuses out in the wake of the star, but then reverses motion and trails the star moving in a flow that resembles accretion. Shocks are apparent at this wake in the ideal EoS that increases the diffusing out of material from the star.

In Fig. 9, we plot only the salient quantities that demonstrate the level of accuracy of the simulations. In panel (a), we demonstrate how the simulation time is approximately the time of the fundamental period of pulsation and the amplitude is similar to the stationary star simulations of Fig. 7(a). The error on the star’s velocity is less than 10^{-3} for *all* simulations (Fig. 9b), where OCTO-TIGER simulations have an error that is about 10 times smaller than the OCTO-TIGER simulations in the case of the slow moving star of Fig. 8(b) (Section 3.6.1). In Fig. 9(c), we see that the ambient medium momentum in OCTO-TIGER simulations remains negligible. After an oscillatory period, we expect it to converge to a value that is of the order of 10^{-9} , which is over 10 times smaller than for the equivalent, slow moving test. Finally, in Fig. 8(d) the mass in the original polytrope is also retained with a precision that is 10 times that of the slow moving polytrope. An amount that is relatively high if we consider the short simulation time.

Click on the link³ to view a movie of the intermediate velocity translating star.

3.6.3 Star translating at a high velocity

For the highest velocity simulation, at Mach 5.2 with respect to the sound speed of the ambient medium, the running time is shorter than the dynamical time of the star so the star does not relax. There is a slight overall mass increase in the two unigrid simulations, but not for the AMR one (Fig. 10a) between the first and fifth time-steps in the simulation. This is due to more mass leaving a cell than the mass in that cell at the sharp stellar edges of the polytrope. Since the mass in a cell is not allowed to be zero nor negative, mass is added to reach the density floor value, hence introducing mass in the grid. At subsequent time-steps when numerical diffusion has softened all edges, this does not happen. This is not a large problem and can be resolved with a reduction of the Courant time to values ~ 0.14 (see justification in Section 2.7).

³<https://youtu.be/b3MMyDCPY60>

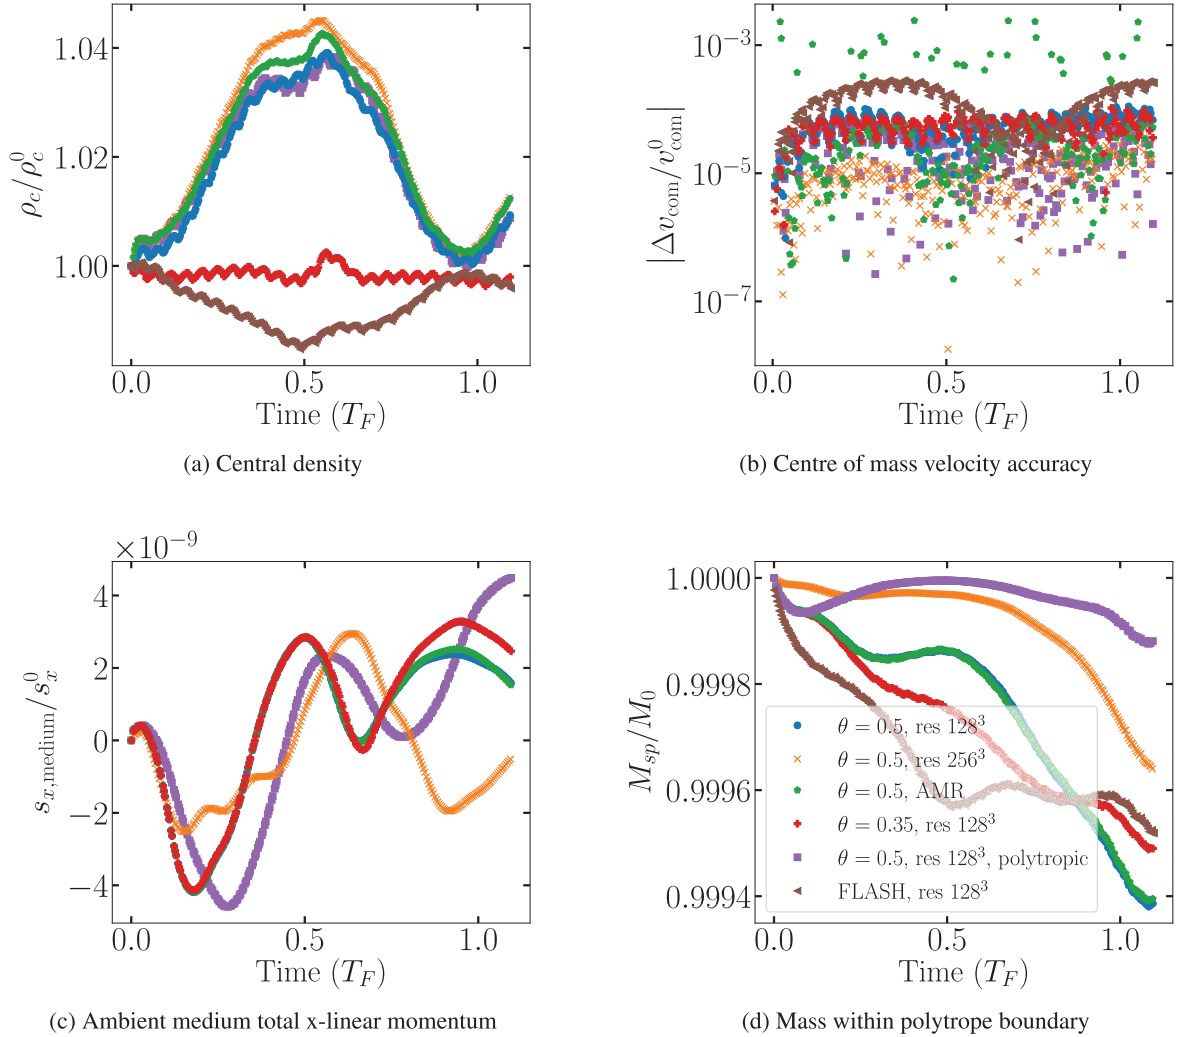


Figure 9. The translating star benchmark test with the star moving at Mach 5.2×10^{-2} (with respect to the diffuse medium). We plot various physical quantities of interest over time. Time is in units of T_F , the fundamental pulsation period.

In Figs 10(b)–(d), we see that the velocity error, ambient x -momentum, and mass in sphere retain their values to a precision that is comparable to that of the polytrope moving at intermediate speed (Section 3.6.2). The best simulation is the high-resolution unigrid with $\theta = 0.5$ (yellow line), something that is also true but to a lesser extent for the other two velocity tests.

Click on the link⁴ to view a movie of the high velocity translating star.

3.7 Star rotating in the grid

Here, we simulate a rotating star in equilibrium with OCTO-TIGER and FLASH. Its rotation profiles, $\Omega(r)$, should remain flat whether the simulation is run in the inertial frame or in a rotating frame of reference, over a certain amount of time.

We constructed an equilibrium profile of a rotating polytrope using the self-consistent field (SCF) method that will be described in Section 4.1. We then interpolated the profile in the OCTO-TIGER and

FLASH grids using the same method. The fast-spinning star has an oblate shape with a ratio of 2/3 between polar and equatorial radii. The angular velocity is 0.52 in code units (0.0023 and 0.19 rad s^{-1} for the MS and the WD models, respectively), and the rotational period is ≈ 12 code units (45 min and 32.3 s for the MS and the WD models, respectively).

We carry out simulations in the inertial frame (both with OCTO-TIGER and FLASH) and in the rotating frame (only with OCTO-TIGER). For each of these frames of reference, we carry out six simulations with the same combinations of resolution, θ parameter, and EoS as done for the stationary star in Section 3.5.

In Figs 11(a) and (b), we plot the familiar evolution of the core density scaled to the initial value. With an exception of the FLASH simulation that has a declining trend, reaching 85 per cent of its initial values over 25 rotation periods, all the OCTO-TIGER simulations retain the initial core density value within 5 per cent with no noticeable difference between rotating and inertial frame. In Fig. 11(c), we plot the diffusion of the stellar mass out of the original boundary in the inertial frame. Once again FLASH performs worse with a loss of 0.5 per cent of the mass over 25 rotation periods, but most of

⁴<https://youtu.be/rZiLMJ6Cxc4>

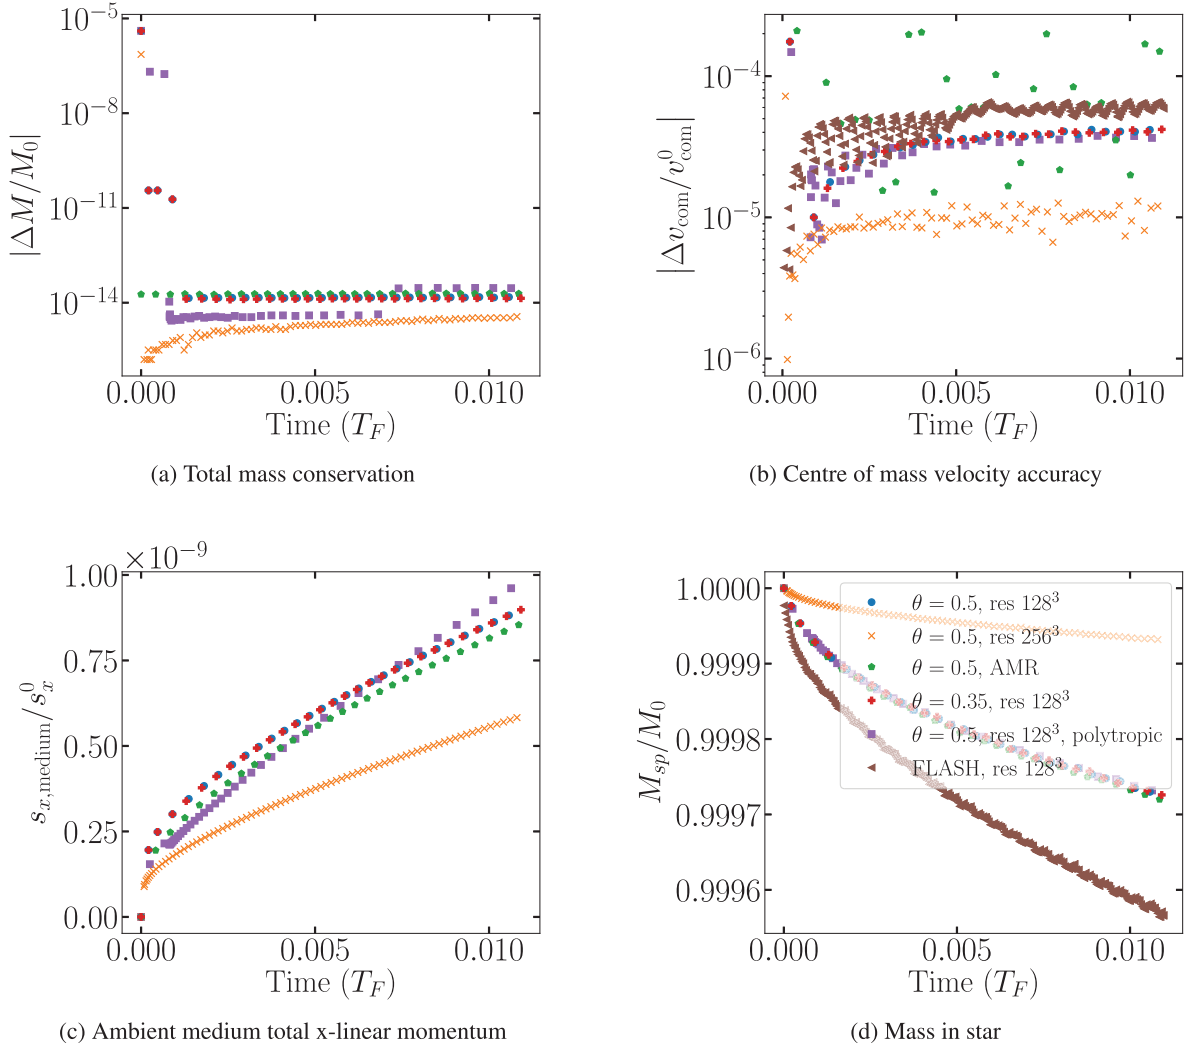


Figure 10. The translating star benchmark test with the star moving at Mach 5.2 (with respect to the diffuse medium). We plot various physical quantities of interest over time. Time is in units of T_F , the fundamental pulsation period.

the OCTO-TIGER profiles are a close second with the polytropic EoS simulation performing distinctly better. Comparing this with panel (d) we see how all OCTO-TIGER simulations perform much better in the rotating frame, losing at most 0.05 per cent of the mass. Similarly, Fig. 11(e) shows that in the inertial frame, the high-density inner star diffuses out moving to twice its original radius. The polytropic EoS simulation is by far the best, with FLASH also performing well. In this case the rotating frame, panel (f), does not perform a great deal better than the inertial frame.

The acid test is the evolution of the rotation profile. All simulations in the inertial frame (Fig. 11g) do not retain a flat profile (with the notable exception of the polytropic EoS one), while those in the rotating frame (Fig. 11h) systematically do. However, in the inertial frame, by the end of the simulations, the spin of the gas that has diffused out of the original volume is distinctly less for all simulations than in the rotating frame. The FLASH simulation (in the inertial frame) does not retain a flat profile, but it has the least diffusion of gas out of the boundary of the sphere of all simulations.

Finally, in Fig. 12 we show the degree of conservation we achieve in the OCTO-TIGER simulations. This highlights the importance of

including the angular momentum correction (AMC) of the gravity solver for the conservation of energy (upper panels) as outlined in Section 3. In the lower panels, we plot the deviation from angular momentum conservation $\Delta J_z/J_z^0$, where $\Delta J_z = J_z(t) - J_z^0$, and $J_z(t) = \int_V (x s_y - y s_x) dV$ and J_z^0 are the z -angular momentum inside the simulation domain, and the initial z -angular momentum, respectively, and where s_x and s_y are the x and y components of the inertial momenta, respectively. As we do not take into account outflows, we can compare between OCTO-TIGER and FLASH, showing OCTO-TIGER outperforms FLASH in ideal gas simulations. Additionally, by comparing between simulations with outflow (blue circles) and reflect (grey downward pointing triangles) boundary conditions, we find that outflows only very slightly affect angular momentum conservation. Overall, the AMR simulation conserves angular momentum the best, and simulations that include AMC conserve angular momentum better than the simulation without this correction.

In conclusion, this test gives an excellent idea of the type of numerical diffusion we can expect. The rotating frame outperforms the inertial frame, and the polytropic EoS seems to be doing best

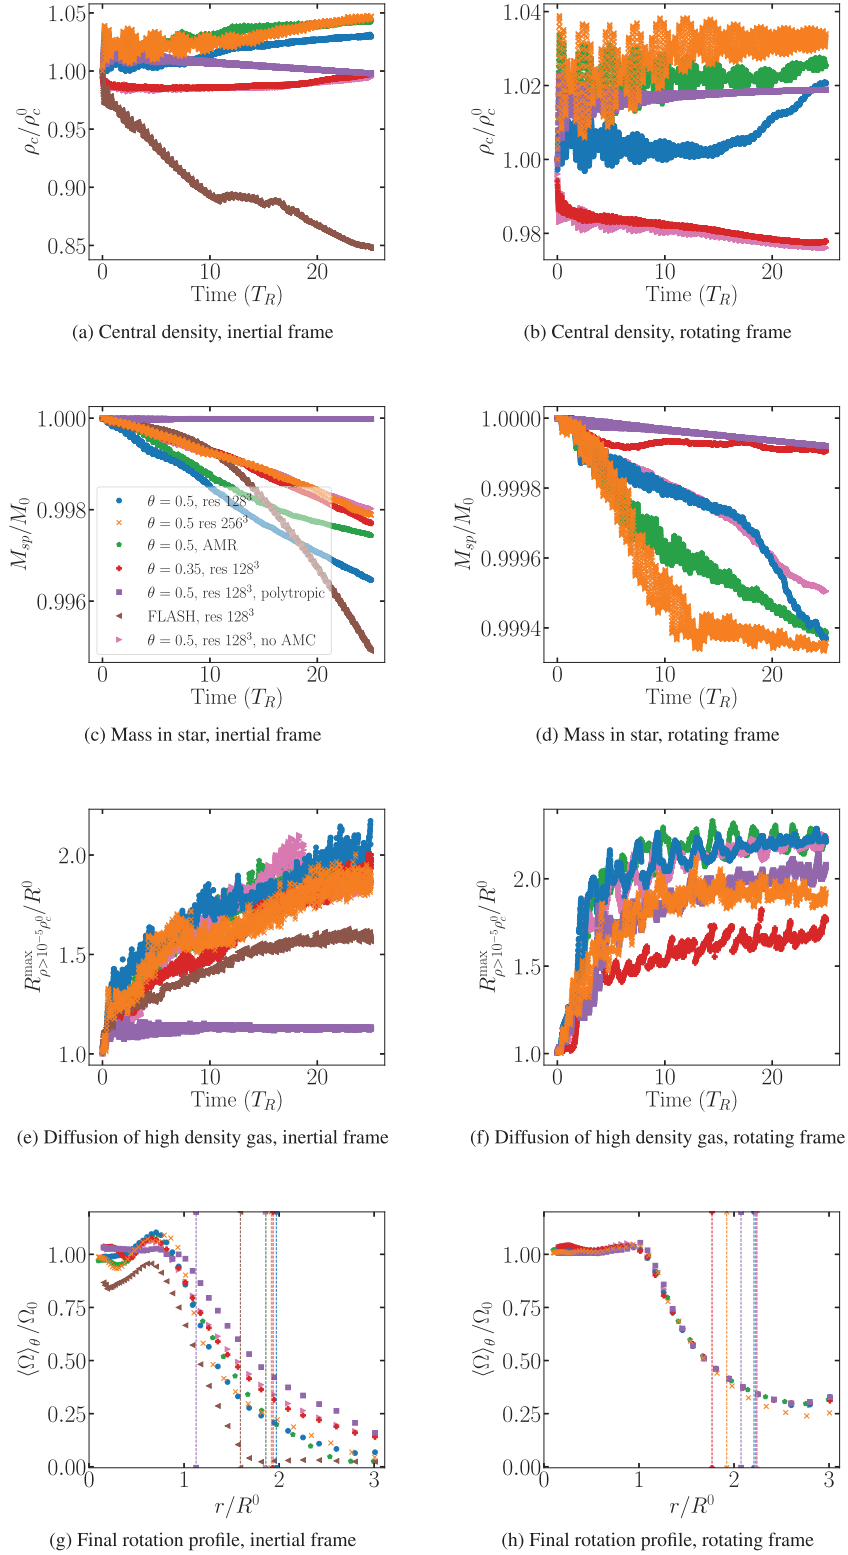


Figure 11. The rotating star test with $\Omega_0 = 0.52$ code units (0.0023 and 0.19 rad s^{-1} for the MS and the WD models, respectively). Several quantities are compared. The vertical lines in panels (g) and (h) are the maximum distance of high-density gas at the end of the simulation (rightmost values of the curves in panels e and f, respectively).

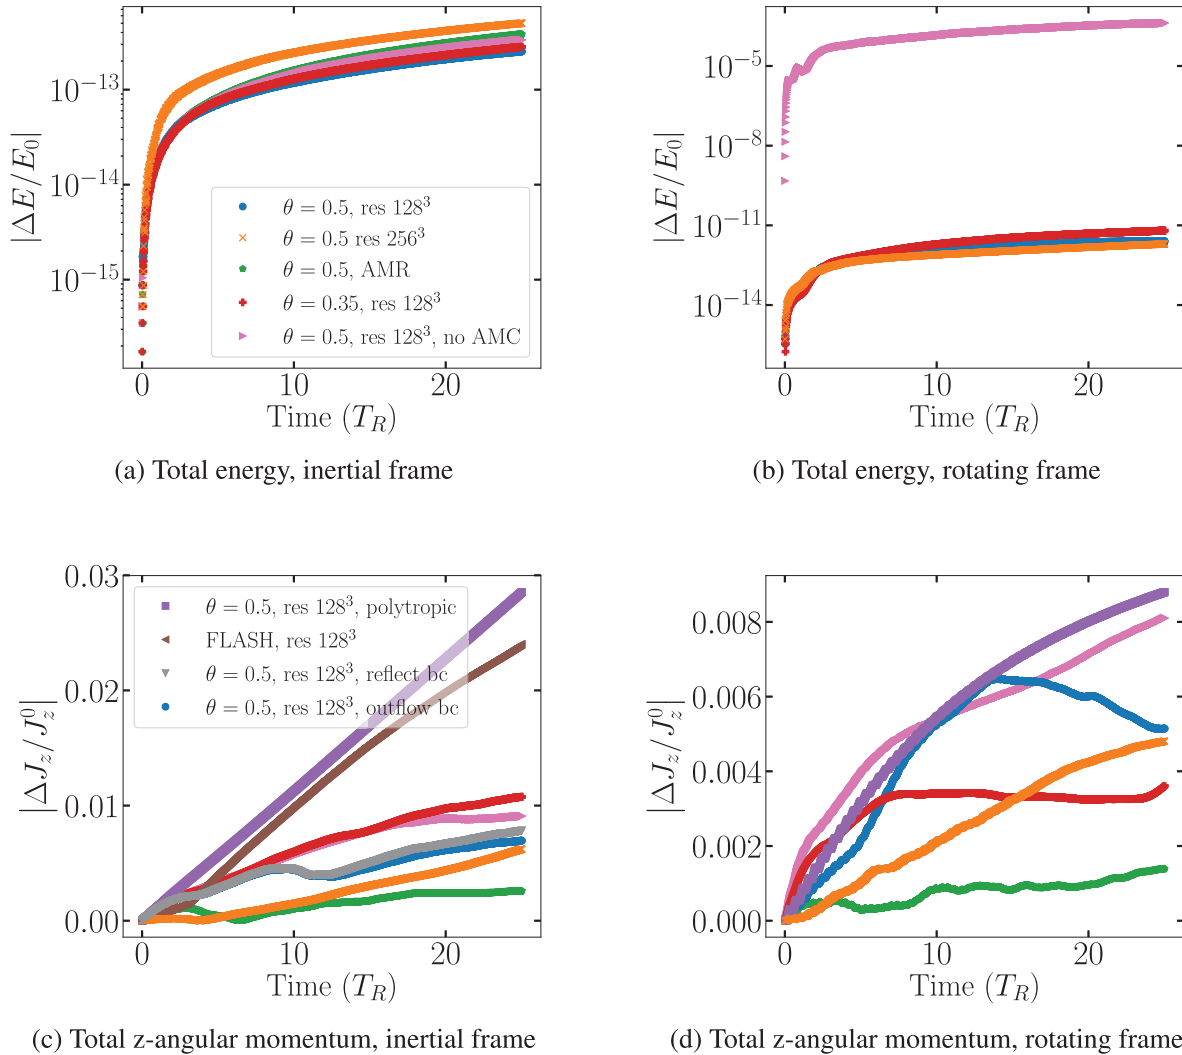


Figure 12. The rotating star test with $\Omega_0 = 0.52$ code units (0.0023 and 0.19 rad/s for the MS and the WD models, respectively). Conservation of energy and angular momentum. All simulations used an ideal gas EoS except the simulation shown in purple squares, which used a polytropic EoS.

in almost all cases. FLASH performs the worst except in retaining a stable high-density sphere inside the star and it can also be seen that a higher gravity accuracy ($\theta = 0.35$) improves the performance.

Click on the link⁵ to view a movie of the rotating star in the rotating frame.

4 BINARY MERGER BETWEEN TWO POLYTROPES WITH A MASS RATIO OF 0.5

Here, we present the results of two binary merger simulations performed in the rotating frame, starting from identical initial conditions, but differing in the adopted EoS. The donor star is half the mass of the accretor and initially fills its Roche lobe. Both stellar spins are synchronized to the orbital frequency. In the first simulation, the two stars are constructed using the same polytropic EoS ($\epsilon_1 = \epsilon_2 = 1$), with polytropic index $n = \frac{3}{2}$. In the second simulation, we use an

ideal gas EoS (with dual energy parameters $\epsilon_1 = 0.001$ and $\epsilon_2 = 0.1$ from equations 11 and 13).

Dominic check all numbers pertaining resolution, grid, and scaling: We are motivated to model such a system for two reasons. First, by appropriate choice of units, this system approximates a binary consisting of two low-mass WDs. We have set the unit conversion factors so that the accretor has a mass of $0.6 M_\odot$, the donor a mass of $0.3 M_\odot$ (1 code mass unit is 1 solar mass). The computational domain is 40 code length units on a side, or 1.7×10^{11} cm (1 length units is 4.36×10^9 cm). With this scaling, the separation between the two stars is $6 \times 10^{-2} R_\odot$, and the period is 2.7 min. The second reason is that a similar system was modelled by Motl et al. (2017) providing us with a comparison simulation as a further verification of OCTO-TIGER.

We use nine levels of refinement, with the highest level having a grid spacing of 4.2×10^7 cm. The accretor is 50 cells across along its longest axis (the spinning stars have slightly larger equatorial than polar radii) and the donor is 79 cells.

Following Motl et al. (2017), we initially drive the stars into deep contact by systematically removing angular momentum. This

⁵https://youtu.be/7ZrxGJW2J_Y

is accomplished by adding source terms to the evolution equations for the x and y components of the momentum (equation 2),

$$\begin{aligned} S_{S_x, \text{driving}} &= -\frac{y}{x^2 + y^2} (xu_y - yu_x) f_{\text{driving}} \\ S_{S_y, \text{driving}} &= +\frac{x}{x^2 + y^2} (xu_y - yu_x) f_{\text{driving}} \end{aligned} \quad (46)$$

where x and y are the x and y components of the position vector \mathbf{x} , u_x and u_y are the x and y component of the inertial frame velocity, \mathbf{u} , and f_{driving} is the driving rate in units of inverse time. This has the effect of reducing the z -angular momentum component of a given cell at a logarithmic rate of f_{driving} . It does not affect the cylindrical radial or z linear momenta. Similarly to what was done by Motl et al. (2017), we drive the system together at a rate of 1 per cent per orbit for the first 2.7 orbits.

4.1 The initial stellar model

The initial conditions for our binary simulations were produced using our own variant of the SCF method (Hachisu 1986a,b; Even & Tohline 2009; Kadam et al. 2016). Given a barotropic EoS, $p = p(\rho)$, the SCF method allows one to construct an equilibrium model of a binary system with synchronously rotating stars. OCTO-TIGER's SCF can be used to construct systems with a polytropic or bi-polytropic EoS or a cold WD EoS. The user selects the masses of each star and the Roche lobe filling factor for the donor. When each star has a different EoS, the Roche lobe filling factor for the accretor must also be specified. Here, we document OCTO-TIGER's SCF as it pertains to the construction of the initial conditions for the binary models in this paper. These models use the same structural polytropic EoS for each star. One of the models is evolved with the same polytropic EoS, and the other with an ideal gas EoS.

The effective potential is defined as

$$\phi_{\text{eff}} = \phi - \frac{1}{2} \Omega_{\text{orb}} r^2, \quad (47)$$

where ϕ is the gravitational potential, Ω_{orb} is the rotational frequency of the binary, and r is the distance to the axis of rotation. Using the effective potential, the hydrostatic equilibrium equation in the rotating frame can be written

$$\frac{1}{\rho} p + \phi_{\text{eff}} = 0. \quad (48)$$

For the EoS, we use:

$$p_{\text{poly}} = K \rho^{1+\frac{1}{n}}, \quad (49)$$

where K is the polytropic constant and n is the polytropic index. Combining equations (48) and (49) and integrating we arrive at:

$$K (1+n) \rho^{\frac{1}{n}} + \phi_{\text{eff}} = C_{1|2}, \quad (50)$$

where the constant on the RHS is either C_1 for the primary, accreting star or C_2 for the secondary, donor star. While in general the polytropic constant, K , may be different for each star, in this paper, they are the same.

The SCF solves for the initial conditions iteratively. The user selects the mass of each star, the polytropic index, the initial separation, and the Roche filling factor of the donor. (Note that while the user specifies the initial separation, it is the orbital angular momentum that is held constant.) The Roche filling factor is defined as

$$f = \frac{\phi_{\text{edge}} - \phi_C}{\phi_{L1} - \phi_C}, \quad (51)$$

where ϕ_{edge} is the effective potential at the edge of the donor, ϕ_C is the effective potential at the centre of mass of the donor, and ϕ_{L1} is the effective potential at L_1 , the first Lagrange point.

For the initial iteration, we place two polytropic stars in the grid at the desired separation and small enough that they are within their respective Roche lobes. The gravity solver is also called before the iterations begin. For each iteration, we: (1) multiply the densities of each star by a constant factor for each star, such that the result yields the desired mass of each star; (2) advect the entire grid such that the centre of mass lies at the centre of coordinates; (3) set the new value for Ω_{orb} using

$$\Omega_{\text{orb}} \rightarrow \frac{J_{0,\text{orb}}(M_1 + M_2)}{M_1 M_2 a^2}, \quad (52)$$

where $J_{0,\text{orb}}$ is the initial orbital angular momentum, M_1 and M_2 are the fixed masses of the accretor and donor, and a is the orbital separation for the current iteration; (4) compute the integration constants C_1 and C_2 ,

$$C_1 = K (1+n) \rho_1^{\frac{1}{n}} + \phi_1, \quad (53)$$

and

$$C_2 = (1-f) \phi_C + f \phi_{L1}, \quad (54)$$

where ρ_1 is the maximum density of the accretor and ϕ_1 is the effective potential at the centre of the accretor; (5) compute a new value for K ,

$$K = \frac{C_2 \phi_2}{(n+1) \rho_2^{\frac{1}{n}}}, \quad (55)$$

where ρ_2 is the density at the centre of the donor, and ϕ_2 is the effective potential at the centre of the donor ($\phi_2 = \phi_C$); (6) compute a new density value at each point on the grid using

$$\rho = \left(\frac{\phi_{\text{eff}} - C_{1|2}}{K (n+1)} \right)^n, \quad (56)$$

for the accretor or donor region as appropriate; and (7) solve for the potential and repeat the process from step (1).

4.2 Simulation results

In each model, mass transfer begins within the first couple of orbits. The transfer rate grows and the donor is tidally disrupted at approximately $21.2P_0$ for the polytropic and $19.6P_0$ for the ideal gas model. Here, P_0 is the initial orbital period, which is the same for each model. The merged remnant is then evolved for a couple of additional P_0 .

In Fig. 13, we show various images at $t = 11.7P_0$ for the polytropic simulation and $t = 12.6P_0$ for the ideal gas simulation. Panels (a) and (b) show density slices in the equatorial plane for the two EoS choices. OCTO-TIGER is able to track the material that was originally part of the donor, and the donor mass fraction is displayed in panels (c) and (d). The polytropic model clearly exhibits the same polygonal resonances at high mass transfer rates that were first reported in D'Souza et al. (2006) and were later confirmed in the SPH simulations reported in Motl et al. (2017). As in Motl et al. (2017), these resonance patterns do not appear in the ideal gas model because the accretion structure is thicker due to shock heating. In Figs 13(e) and (f), we show a density cut at the same times, showing that while the ideal gas EoS simulation produces L_2 and L_3 outflow, the polytropic EoS simulation only shows a hint of L_2 outflow by that time. The level of detail clearly visible near L_1 is consistent with the Roche geometry.

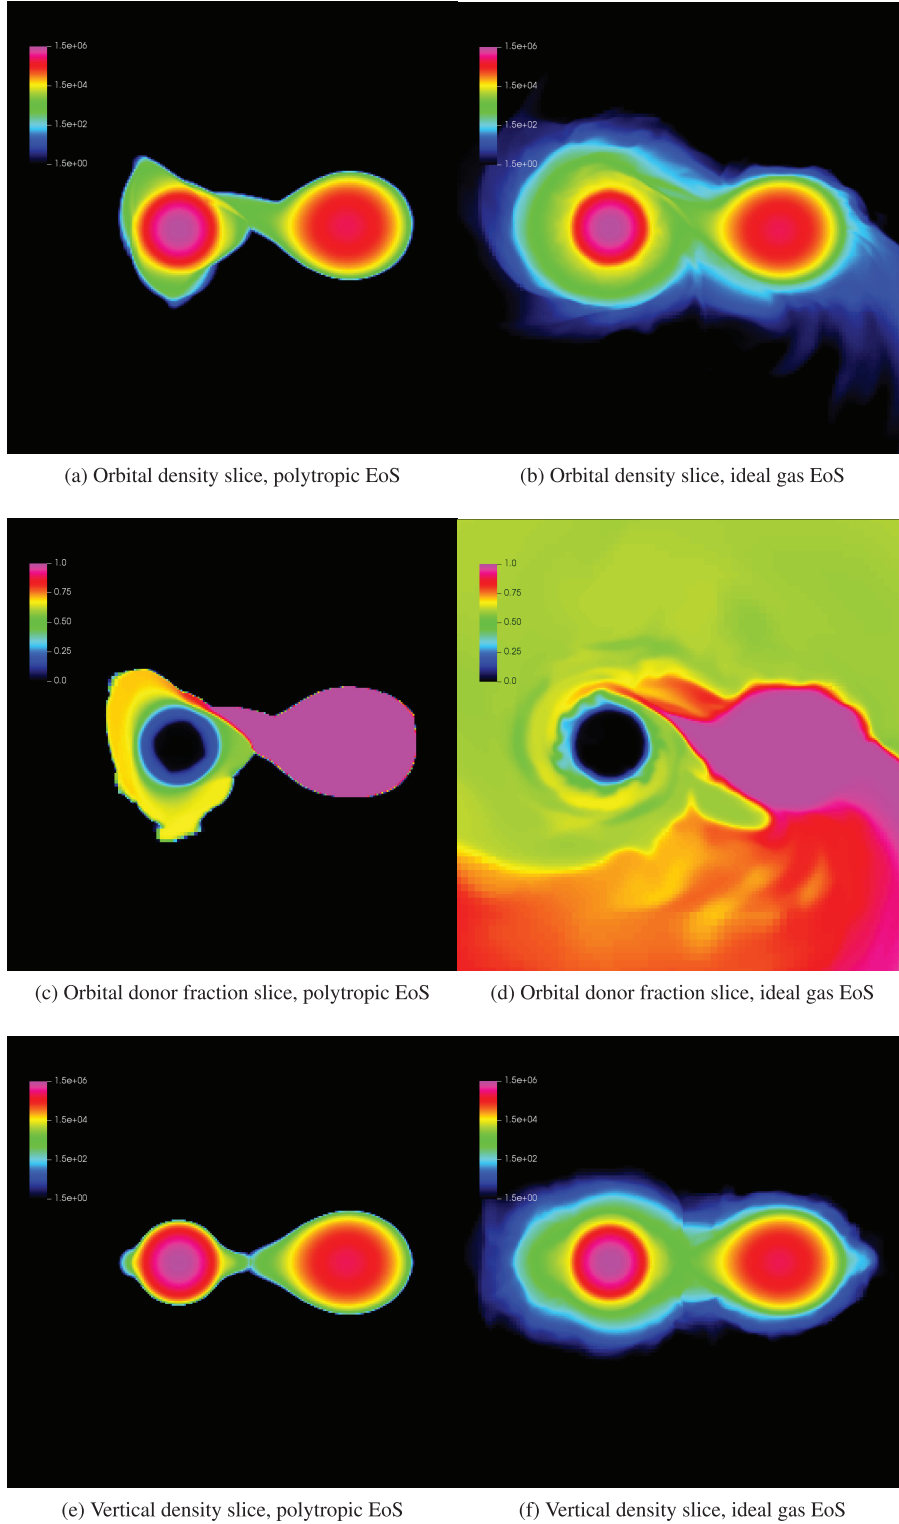


Figure 13. Snapshots from our $q = 0.5$ simulations with a polytropic EoS (left column, at time $11.7 \times P_0$, where P_0 is the initial orbital period) or an ideal gas simulation (right column, at time $12.6 \times P_0$). The box size is 11×10^9 cm in all cases. The density scale is logarithmic and runs from 1.5 to 1.5×10^6 g cm^{-3} .

Figs 14(a) and (b) are from earlier in the evolution at approximately $t = 7P_0$ in each model. The resonance patterns in the polytropic model progress from pentagonal shaped, to box shaped, to triangular shaped (clearly seen in in Fig. 14a). Figs 14(c) and (d) show the system as tidal disruption is occurring. As the donor is disrupted, significant

mass flows through the L_2 Lagrange point, resulting in the ‘tails’ seen in the figure.

The images in Fig. 15 are taken from two orbits after tidal disruption and merger has occurred. Panels (a) and (b) show equatorial density slices, while panels (c) and (d) show the donor fraction. The

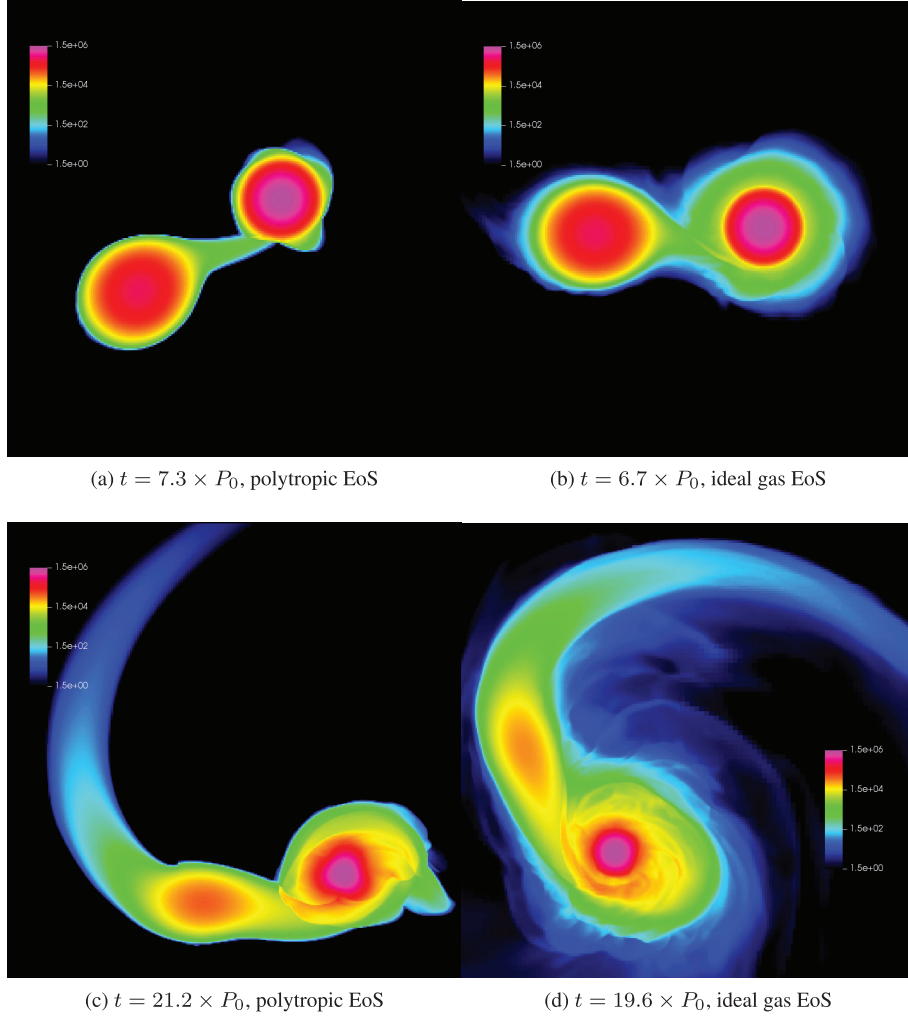


Figure 14. Density slices on the orbital plane for our $q = 0.5$ simulations with a polytropic EoS (left column) or an ideal gas EoS (right column). The box size is 11×10^9 cm (top row) or 16×10^9 cm (bottom row). The density scale is logarithmic and runs from 1.5 to 1.5×10^6 g cm $^{-3}$. The snapshots are taken at times 7.3 (a), 6.7 (b), 21.2 (c), and $19.6 \times P_0$ (d). See Table 1 for description.

triangular shaped resonance pattern is still evident near the core in the polytropic model. In the ideal gas model, there is still a remnant of the donor above and just to the left of centre. The bottom images are shown with a larger spatial and density range to reveal the large-scale low-density structures that form from the merger.

4.2.1 Simulation diagnostics

In order to measure the simulated properties of each star in the binary, we must first establish how to locate the stellar surface of each star. OCTO-TIGER accomplishes this iteratively. For each iteration, OCTO-TIGER updates the centres of mass of each star and the orbital frequency. Summations are done over the grid cells within the regions defined for each star, and indices i, j , and k referring to grid cells are implied on the RHS of the next two equations. We thus define the centres of mass of each star as

$$\mathbf{x}_{1|2} = \frac{\sum^{1|2} \rho \mathbf{x} \Delta x^3}{M_{1|2}}. \quad (57)$$

The velocity of the centre of mass of each star is

$$\mathbf{u}_{1|2} = \frac{\sum^{1|2} \rho \mathbf{u} \Delta x^3}{M_{1|2}}. \quad (58)$$

Given $\mathbf{x}_{1|2}$ and $\mathbf{u}_{1|2}$, we define the orbital angular frequency as

$$\Omega_{\text{orb}} = \frac{[(\mathbf{x}_1 - \mathbf{x}_2) \times (\mathbf{u}_1 - \mathbf{u}_2)] \cdot \mathbf{e}_z}{|\mathbf{x}_1 - \mathbf{x}_2|^2}, \quad (59)$$

where \mathbf{e}_z is the unit vector in the z direction, out of the orbital plane.

For the first iteration, we compute the centres of mass and orbital frequency based on an initial guess for the bounding surface of each star. OCTO-TIGER evolves the original accretor and original donor mass densities as passive scalars, allowing it to track donor material within the accretor. We use the volumes defined by these densities to seed the first iteration, with cells containing majority accretor (donor) material flagged for the accretor (donor).

With the stellar volumes defined, we can then calculate the centres of mass and orbital frequency. Beginning with the second iteration (out of a total of five), we use the acceleration from the effective potential to find the regions of each star. This acceleration is

$$\mathbf{g}^{\text{eff}} = \mathbf{g} + \mathbf{r} \Omega_{\text{orb}}^2, \quad (60)$$

where \mathbf{r} is the radial vector distance to the axis of rotation. We define the quantities

$$q_{1|2} = \mathbf{g}^{\text{eff}} \cdot \frac{\mathbf{x} - \mathbf{x}_{1|2}}{|\mathbf{x} - \mathbf{x}_{1|2}|}. \quad (61)$$

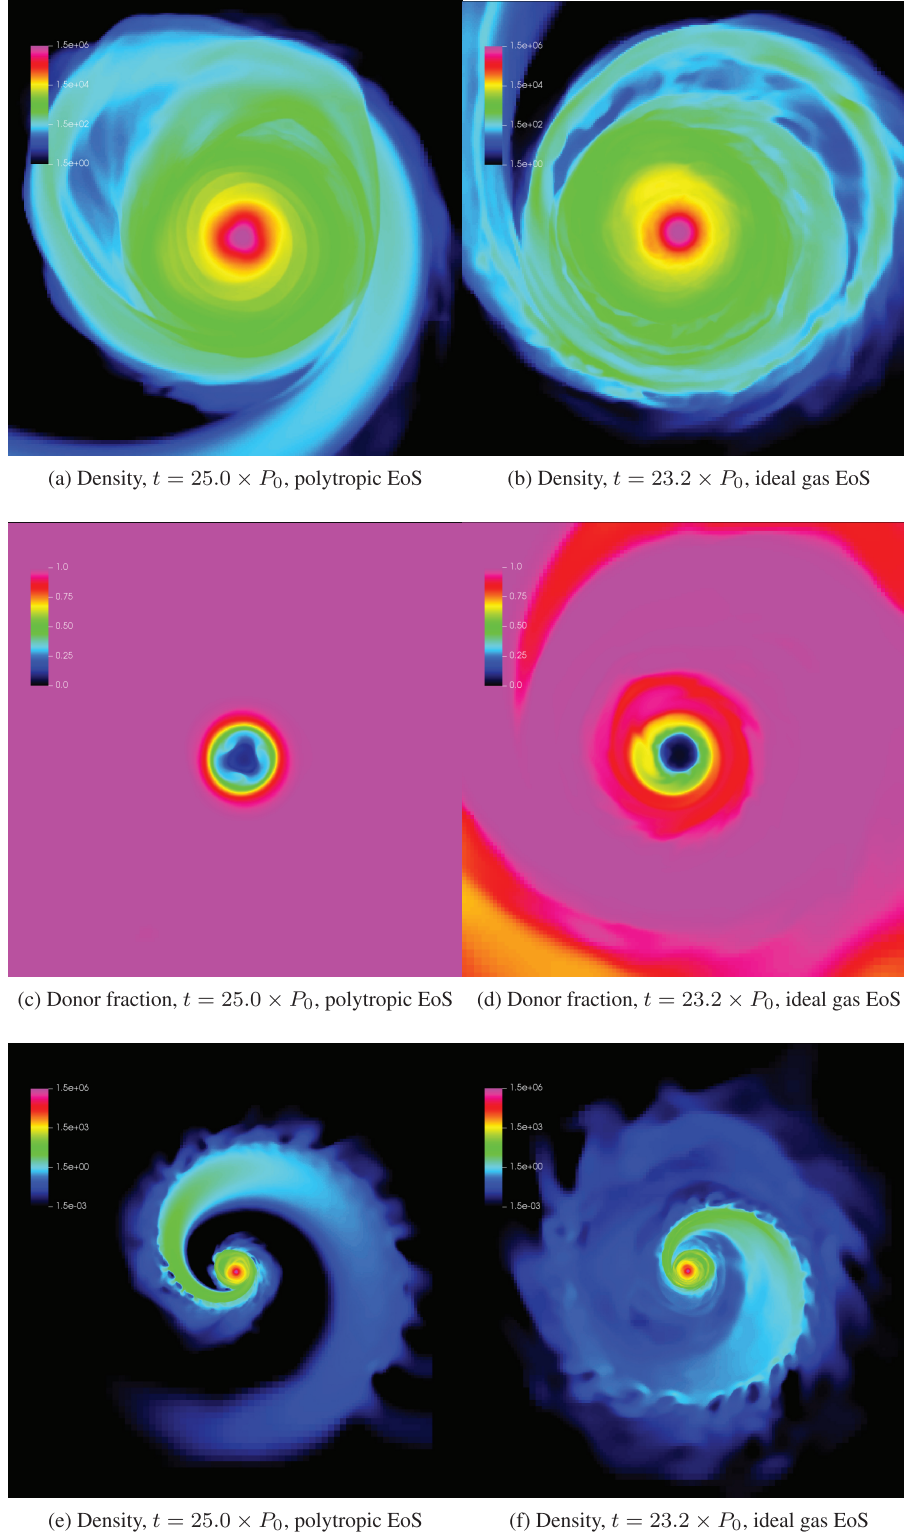


Figure 15. Slices on the orbital plane for our $q = 0.5$ simulations with a polytropic EoS (left column) or an ideal gas EoS (right column). The box size is 20×10^9 cm (top two rows; top row: logarithmic density scale with range $1.5\text{--}1.6 \times 10^6$ g cm $^{-3}$) or 200×10^9 cm (bottom row; logarithmic density scale with range 1.5×10^{-3} to 1.6×10^6 g cm $^{-3}$). The snapshots are taken at times 25.0 (a, c, e), and $23.2 \times P_0$ (b, d, f). See Table 1 for description.

If $\min(q_1, 0) < \min(q_2, 0)$, the cell belongs to the accretor. If $\min(q_1, 0) > \min(q_2, 0)$, the cell belongs to the donor. For the case where $q_1 \geq 0$ and $q_2 \geq 0$, the cell is in neither star. This definition by itself may miss some cells in the centres of each star where the

effective gravitational force is not guaranteed to point towards the centre of mass. For this reason, any material within a certain critical radius of the centre of mass of a star is included as part of that star regardless of the values of $q_{1|2}$. This critical radius is $\frac{1}{4}R_{L,1|2}$, where

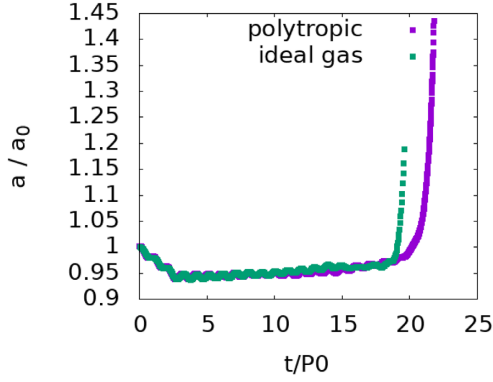


Figure 16. Diagnostic plots from the $q = 0.5$ polytropic and ideal gas simulations. Orbital separation is given as a function of the initial value.

$R_{L,1|2}$ is the Roche lobe radius of each star taken from the point mass approximation.

With the stellar volumes defined, we can measure the properties of each star. The mass of each star is thus defined as

$$M_{1|2} = \sum_{1|2} \rho \Delta x^3. \quad (62)$$

The orbital separation is

$$a = |\mathbf{x}_2 - \mathbf{x}_1|. \quad (63)$$

We show the separation of each model in Fig. 16. Once the driving phase is over, the orbital separation of the two simulations increases with time up until the final merger. This shows the donor does not plunge into the accretor but instead is tidally disrupted (e.g. Figs 14c and d). The matter from the disrupted donor then falls on to the accretor. Our diagnostic plots that measure properties of the individual stars do not extend to this phase because there is no longer a clearly defined donor.

Computing the time rate of change of the two stars' masses, $\dot{M}_{1|2}$, requires filtering out high frequency noise, here defined as anything greater than $(\frac{1}{2}\Omega_{\text{orb}})$. The regions defining each star change with time and can vary slightly from one time-step to the next. The high frequency part of this signal needs to be filtered out to measure a meaningful $\dot{M}_{1|2}$. We apply a windowed sinc filter using an exact Blackman window with frequencies of $\frac{1}{2}\Omega_{\text{orb},0}$ and $\frac{3}{20}\Omega_{\text{orb},0}$. We apply this filter to $M_{1|2}$ and then compute the time-derivatives with nearest neighbours by a first-order discrete scheme.

In Fig. 17 (top left panel), we show the donor mass-loss rate, normalized to the donor's mass. The time derivative is in units of initial orbital periods. The donor initially loses less than 1 per cent per orbit, increasing to ~ 1 per cent, and then ramping up towards the time when the donor is tidally shredded. In Fig. 17 (top right panel), we show the amount of mass that is no longer a part of either star. The two star system loses only a small fraction of the donor's mass throughout most of the evolution. This can be mass unbound from the system or mass in a common envelope, but still bound to the system. The ideal gas model loses mass at a higher rate. This is because the accreted gas is shock heated and therefore has higher pressure. It builds up a thicker layer around the accretor, and some of this gas is able to escape.

In Fig. 17 (bottom panels), we plot the mass that is outside the surface of either stars, and the mass that is unbound from the system. Unbound mass is that for which the inertial frame kinetic energy density exceeds the inertial frame potential energy. Further from the system centre of mass, where most of the unbound material is likely

to be located, this indicator is more reliable than using rotating frame quantities. Most of the mass remains bound. In the case of the ideal gas model, about 10 per cent of the mass leaves the stellar boundaries, but only about 1 per cent of the mass actually becomes unbound from the system. In the case of the polytropic model, these figures are even lower due to lack of shock heating.

The orbital angular momentum is

$$J_{\text{orb}} = \frac{M_1 M_2}{M_1 + M_2} \Omega_{\text{orb}} a^2. \quad (64)$$

We show J_{orb} for both models in the top left panel of Fig. 18 as a function of the initial orbital angular momentum ($1.5 \times 10^{51} \text{ g m}^2 \text{ s}^{-1}$). As expected, both simulations lose orbital angular momentum throughout their evolution.

Each star has a spin angular momentum

$$J_{1|2} = \sum_{1|2} \rho (\mathbf{x} - \mathbf{x}_{1|2}) \times (\mathbf{u} - \mathbf{u}_{1|2}) \Delta x^3, \quad (65)$$

and an angular frequency assuming rigid body rotation

$$\Omega_{1|2} = \frac{J_{1|2}}{\sum_{1|2} \rho |(\mathbf{R} - \mathbf{R}_{1|2})|^2 \Delta x^3}, \quad (66)$$

where $\mathbf{R}_{1|2}$ is the projection of $\mathbf{x}_{1|2}$ on the xy -plane. We show $J_{1|2}$ in the middle right and lower right panels of Fig. 18. In each model, the accretor gains angular momentum from the donor. We show $\Omega_{1|2}$ along with Ω_{orb} in Fig. 18.

The gravitational torques exerted on each star are

$$\mathbf{T}_{1|2} = \sum_{1|2} (\mathbf{x} - \mathbf{x}_{1|2}) \times \rho \mathbf{g}. \quad (67)$$

The gravitational torque exerted on the orbit is

$$\mathbf{T}_{\text{orb}} = \mathbf{x}_1 \times \sum \rho \mathbf{g} + \mathbf{x}_2 \times \sum \rho \mathbf{g}. \quad (68)$$

The z -component of \mathbf{T}_{orb} plays an important role in whether or not the binary survives mass transfer. This represents the rate at which the tidal interaction can restore the orbital angular momentum lost to mass transfer (Figs 19a–c).

D'Souza et al. (2006) give an expression relating the time rate of change of the separation to systematic sources of orbital angular momentum, gravitational torque, and the mass transfer rate. This is

$$\frac{\dot{a}}{2a} = \left(\frac{\dot{J}}{J_{\text{orb}}} \right)_{\text{sys}} + \mathbf{T}_{z,\text{orb}} - \frac{\dot{M}_2}{M_2} (1 - q - \sqrt{(1+q)r_h}). \quad (69)$$

The first term on the RHS, $\left(\frac{\dot{J}}{J_{\text{orb}}} \right)_{\text{sys}}$, is due to systematic changes in orbital angular momentum due to losses through L_2 . In addition, there is a loss of orbital angular momentum during the first 2.7 orbits because of our artificial driving that needs to be accounted for when using this expression. The second term represents the contribution due to the gravitational torque on the orbit. The final term includes, r_h , the effective 'circularization' radius. This is the radius from the accretor's centre of mass, which has the same specific angular momentum as the material at the L_1 Lagrange point, normalized to the orbital separation. While the first part of the final term, $\frac{\dot{M}_2}{M_2} (1 - q)$, accounts for the change in orbital angular momentum due to the transfer of mass from one star to the other, the second part, $\frac{\dot{M}_2}{M_2} (-\sqrt{(1+q)r_h})$, accounts for the orbital angular momentum that is converted to spin angular momentum in the accretor.

Frank et al. (2002) provide an analytical approximation to r_h ,

$$r_{h,\text{analytic}} = (1 + q) (0.500 - 0.227 \log_{10} q)^4. \quad (70)$$

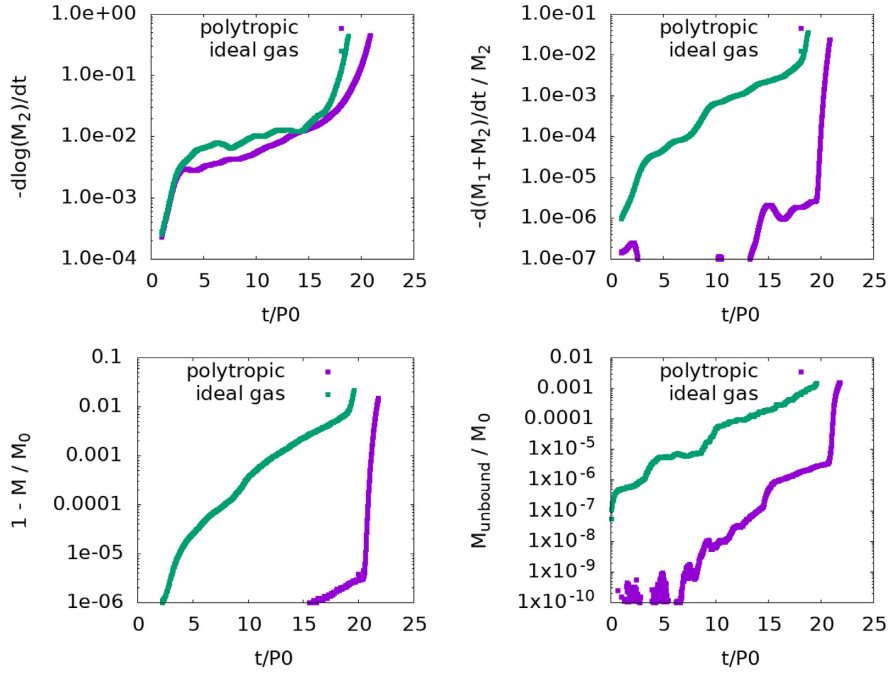


Figure 17. Diagnostic plots from the $q = 0.5$ polytropic and ideal gas simulations. Upper left: donor mass-loss rate; upper right: mass-loss rate from the both stars ($M = M_1 + M_2$, M_0 is the initial mass of the system); lower left: mass outside the boundary of the two stars; lower right: mass unbound from the binary.




Figure 18. Diagnostic plots from the $q = 0.5$ polytropic and ideal gas simulations. Top left: the orbital angular momentum; top middle: the accretor angular momentum; top right: the donor angular momentum. All angular momenta are normalized to the initial value of the orbital angular momentum. Bottom left: the orbital frequency; bottom middle: the spin frequency of the accretor; bottom right: the spin frequency of the donor. All frequencies are normalized to the orbital frequency at $t = 0$.

A more accurate expression that includes the torquing of the stream of accreting gas by the donor is that of Verbunt & Rappaport (1988). However, for the current purpose the expression in equation (70) is sufficiently accurate. The orbital separation, angular momentum driving rate, gravitational torque on the orbit, the mass transfer rate, and mass ratio are all known or can be calculated from the data set, allowing us to compute a value for r_h using equation (69). We show

this quantity compared with $r_{h, \text{analytic}}$ for the polytropic model, in the bottom right panel of Fig. 19. In the beginning the value of r_h varies wildly. This is during the initial period of angular momentum driving. For the majority of the run, r_h is slightly less than $r_{h, \text{analytic}}$, indicating the spin angular momentum actually transferred to the accretor is slightly less than predicted by the analytical theory. Towards the end, the computationally derived r_h begins to exceed the analytical value

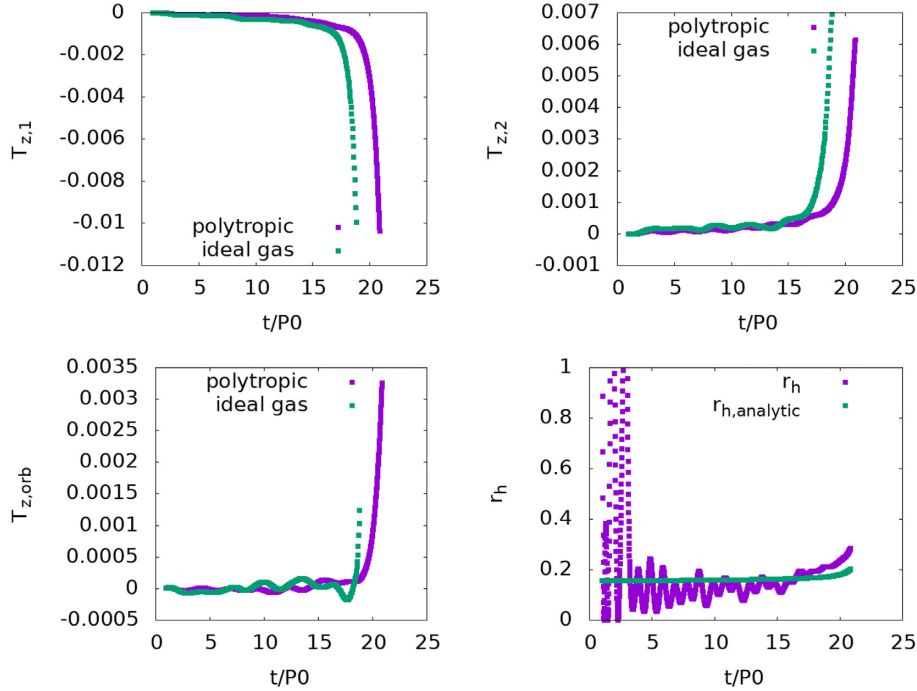


Figure 19. Diagnostic plots from the $q = 0.5$ polytropic and ideal gas simulations. Top left: the gravitational torque on the accretor; top right: the gravitational torque on the donor; bottom left: the gravitational torque on the orbit; bottom right: the circularization radius for the polytropic simulation calculated using equation (69) and for the analytical approximation using equation (70).

$r_{h, \text{analytic}}$. This is because angular momentum is now flowing through the L_2 point, making equation (69) no longer valid.

The internal, kinetic, and potential energies of the stars are

$$E_{I,1|2} = \sum_{1|2} \rho \varepsilon, \quad (71)$$

$$E_{K,1|2} = \sum_{1|2} \frac{1}{2} \rho u^2, \quad (72)$$

and

$$\Phi_{1|2} = \sum_{1|2} \frac{1}{2} \rho \Phi, \quad (73)$$

respectively. Here, the internal energy density, $\rho \varepsilon$, is obtained from the dual energy formalism. We show these energies in Fig. 20. The donor loses internal and kinetic energy to the accretor. Although the ideal gas model includes shock heating, since the energy of the shock-heated material around the accretor is very small compared to the total internal energy of the star, shock heating makes little apparent difference in the plots of total internal energy. As the accretor gains mass, its gravitational well deepens as its potential energy is lowered. The opposite is true of the accretor.

Grid based codes are subject to centre of mass drift, both physical, from the flow of matter out of the grid (which, by the end of the simulations, is 2.0×10^{29} and 1.0×10^{28} g for the ideal gas and polytropic EoS, respectively), and numerical. The latter effect is greatly reduced by using a gravity solver that conserves linear momentum, as OCTO-TIGER does. There are still numerical viscosity effects that can push matter one way or another. In addition, the SCF code does not render initial conditions with the centre of mass perfectly at the coordinate centre. Fig. 21 (top left panel) shows the magnitude of cylindrical radial location of the system's centre of

mass,

$$|\mathbf{R}_{\text{com}}| = \sqrt{x_{\text{com}}^2 + y_{\text{com}}^2} \quad (74)$$

in length units normalized to the finest cell width.

As discussed in Section 2.1, OCTO-TIGER evolves the angular momentum (equation 7) for the purposes of obtaining the error in the angular momentum conservation. The evolved angular momenta, \mathbf{l} , can be compared with the angular momenta obtained from the evolved linear momenta, $\mathbf{x} \times \rho \mathbf{u}$, (equation 2) to obtain the error. We define the global angular momentum conservation violation error, \mathbf{l}_{err} , as

$$\mathbf{l}_{\text{err}} = \sum_{i=1}^v (\mathbf{l}_i - \mathbf{x}_i \times \rho_i \mathbf{u}_i), \quad (75)$$

where the i index refers to a given computational cell and the summation is over the entire computational domain. We show this value over time in Fig. 21 (top right panel). In Fig. 21, bottom left panel, we plot the total angular momentum non-conservation in the computational domain, scaled to the initial total angular momentum. Here, we see that the polytropic EoS conserves momentum to approximately 1 percent over the entire simulation, while for the ideal gas EoS this is considerably worse at 7 percent (the initial drop in angular momentum in the right-hand panel is due to the artificial driving phase and should be ignored). The reason for the difference is that the ideal gas EoS simulation expands more due to shock heating moving gas into the outskirts of the domain where lower resolution dominates with a worse performance of angular momentum conservation. Poorly resolved, sharp momentum density gradients, such as those seen in the outskirts of the simulations (see Fig. 15, bottom panels, where we plot density slices) will always disfavour good conservation.

In Fig. 21, bottom right panel, we show the energy conservation. The value in the y-axis is scaled to the energy after the driving, at




Figure 20. Diagnostic plots from the $q = 0.5$ polytropic and ideal gas simulations. Top left: the accretor's internal energy; top middle: the accretor's kinetic energy; top right: the accretor's potential energy. Bottom left: the donor's internal energy; bottom middle: the donor's kinetic energy; bottom right: the donor's potential energy. All energies are normalized to the absolute value of the total potential energy at $t = 0$.




Figure 21. Diagnostic plots from the $q = 0.5$ polytropic and ideal gas simulations. Top left: the location of the system centre of mass (equation 74) in units of the finest grid cell; top right: the relative angular momentum error (equation 75); bottom left: the total angular momentum loss from the computational domain (the initial decrease corresponds to the angular momentum that is artificially subtracted from the system during the driving phase); bottom right: the total energy conservation error of the system (the values before $x = 2.7$ orbits are meaningless because of the driving).

2.7 orbits. Hence data points before $x = 2.7$ are not meaningful. After that x value, however, we see that the energy conservation is excellent, going from 10^{-14} to 10^{-11} for the ideal gas EoS, or to 10^{-5} for the polytropic EoS.

Without doubt, the polytropic EoS simulation is superior in general, because the ideal gas EoS with more heating results in an inflated gas distribution, a higher rate of mass transfer from donor to accretor, and any quantities that derive from that, such as angular momentum transfer or the speed of the merger are also affected.

Table 5. Computational infrastructure used in this work. QB2 is Louisiana’s Optical Network Infrastructure’s QueenBee2; BR3 is the University of Indiana BigRed3, while Gadi is the Australian National Computational Infrastructure peak machine.

Cluster	CPU	Memory	Interconnect	# nodes	# of cores
QB2	2 × 10 Intel E5-2680v2 Xeon processors (+ 2 GPUs) ^a	64 GB	56 Gb/sec (FDR) InfiniBand	504	10 800
BR3	2 × 12 Intel Xeon Processor E5-2690 v3	64 GB	Cray Aries	930	22 464
Gadi	2 × 24 core Intel Xenon Scalable ‘Cascade Lake’	190 GB	200 Gb/sec (HDR) InfiniBand	3024	145 152

^aThe GPUs are NVIDIA Tesla K20; NVIDIA-SMI 450.51.05, driver version: 450.51.05, CUDA version: 11.0 with 6GB of memory per device.

Table 6. Software dependences of OCTO-TIGER (version 0.8). Note that we used a customized version of Vc. Note that we used the pre-compiled MPI version on the cluster and therefore have some variation there. If the gcc compiler was recent enough to compile HPX and OCTO-TIGER, we used the pre-compiled version on the cluster.

Software	QB2	BR3	Gadi
HPX	1.4	1.4	1.4
VC	1.4.1	1.4.1	1.4.1
BOOST	1.69/1.68	1.68	1.72
HWLOC	1.11.1	1.11.12	1.11.12
JEMALLOC	5.1.0	5.1.0	5.1.0
GCC	8.3/7.4	8.3	8.3
HDF5	1.12/1.8	1.8.12	1.8.12
SILO	4.10.2	4.10.2	4.10.2
CMAKE	3.13	3.13.2	3.16.2
CUDA	10.2	–	–
MPI	MVAPICH2 2.3.2/ OPEN-MPI 4.0	CRAY-MPICH 7.7.10	INTEL-MPI 2019.8.254

These simulations were run on 400 cores over 20 nodes, of QueenBee2 (see Section 5) and took approximately 200 h of wall-clock time, or just over 8 d for a total modest cost of approximately 80 000 CPU-hours. They have between 2 and 8 million cells. If the scaling test in Section 5 are consulted, we would conclude that a doubling of the number of cores would reduce the wall-clock time between 40 and 60 per cent, while any further increase of the number of cores would not be particularly advantageous.

Naturally, this run could be repeated with substantially larger resolution (and larger number of cores), likely leading to an expansion of discovery space. However, increasing resolution with the aim of determining convergence is not straightforward because it is not clear what quantities can be reasonably expected to converge. An increase in resolution automatically lowers the mass transfer rate in the early interaction and lengthens the pre-merger time beyond what can be simulated (see for instance Motl et al. (2017) who compared a 4-million cell WD merger cell simulation to a 47-million cell one, or Reichardt et al. (2019) who compared the early mass transfer before a common envelope interaction between a $\sim 100\,000$ SPH particle simulation to one with 1.3 million particles). What quantities, then, should we expect to converge to a given value with increasing resolution? It is possible that, for example, the increase in mass transfer rate over one orbital period, measured at a specific initial separation may be a quantity that is expected to converge. The determination of how to measure the accuracy of simulation results of this level of complexity is an urgent, on-going area of research.

Click on the link⁶ to view a movie of the $q = 0.5$ ideal gas simulation.

Click on the link⁷ to view a movie of the $q = 0.5$ polytropic EoS simulation.

5 CODE PERFORMANCE AND SCALING

In this section, we present a series of scaling tests to demonstrate the performance of OCTO-TIGER and we make comparisons with equivalent tests carried out with FLASH. In order to compare the two codes, we set the CFL coefficient to be the same in both ($\eta_{\text{CFL}} = 0.4$).

We carried out two types of scaling tests. The first one was performed by executing a number of time-steps (N_{steps}) in the simulation of a stationary star. The second test included a binary system of stars and was designed to stimulate some level of regridding. Tests were performed on three supercomputers, BigRed3 (BR3), QueenBee2 (QB2), and Gadi, with variable resolutions (128^3 , 256^3 , 512^3 , and 1024^3) and variable values of θ (0.35 and 0.5), the gravity resolution parameter. We describe the supercomputer hardware in Table 5. The code and its dependences’ versions used with the three sets of scaling runs are reported in Table 6.

In Figs 22 and 23, we present two measures of code speed: the computational time per time-step and the wall-clock time per time-step. We show this including and excluding the initialization time, regridding time, and the time to write out the output. While the computational time is a feature of the code, the time to write out the output is a feature of the machine and file system used and can therefore vary among otherwise similar computers. These quantities are a measure of the strong scaling properties of OCTO-TIGER. We also show the number of time-steps (N_{steps}) times the number of cells (N_{cells}) processed per second, which measures the amount of work that the cores have available to do. If we had perfect scaling, these curves should be horizontal, indicating the cores have sufficient work, but they tend instead to curve downwards as the addition of more cores results in cores not having sufficient work to do. The reason why the curves tilt downwards is also the introduction of additional overheads, such as message sending, due to the addition of nodes. Note that above a certain high number of cores, each core has computational work of only several subgrids,

⁶<https://www.youtube.com/watch?v=0JD5E7DUImw>

⁷<https://www.youtube.com/watch?v=MfArAQPPHss>

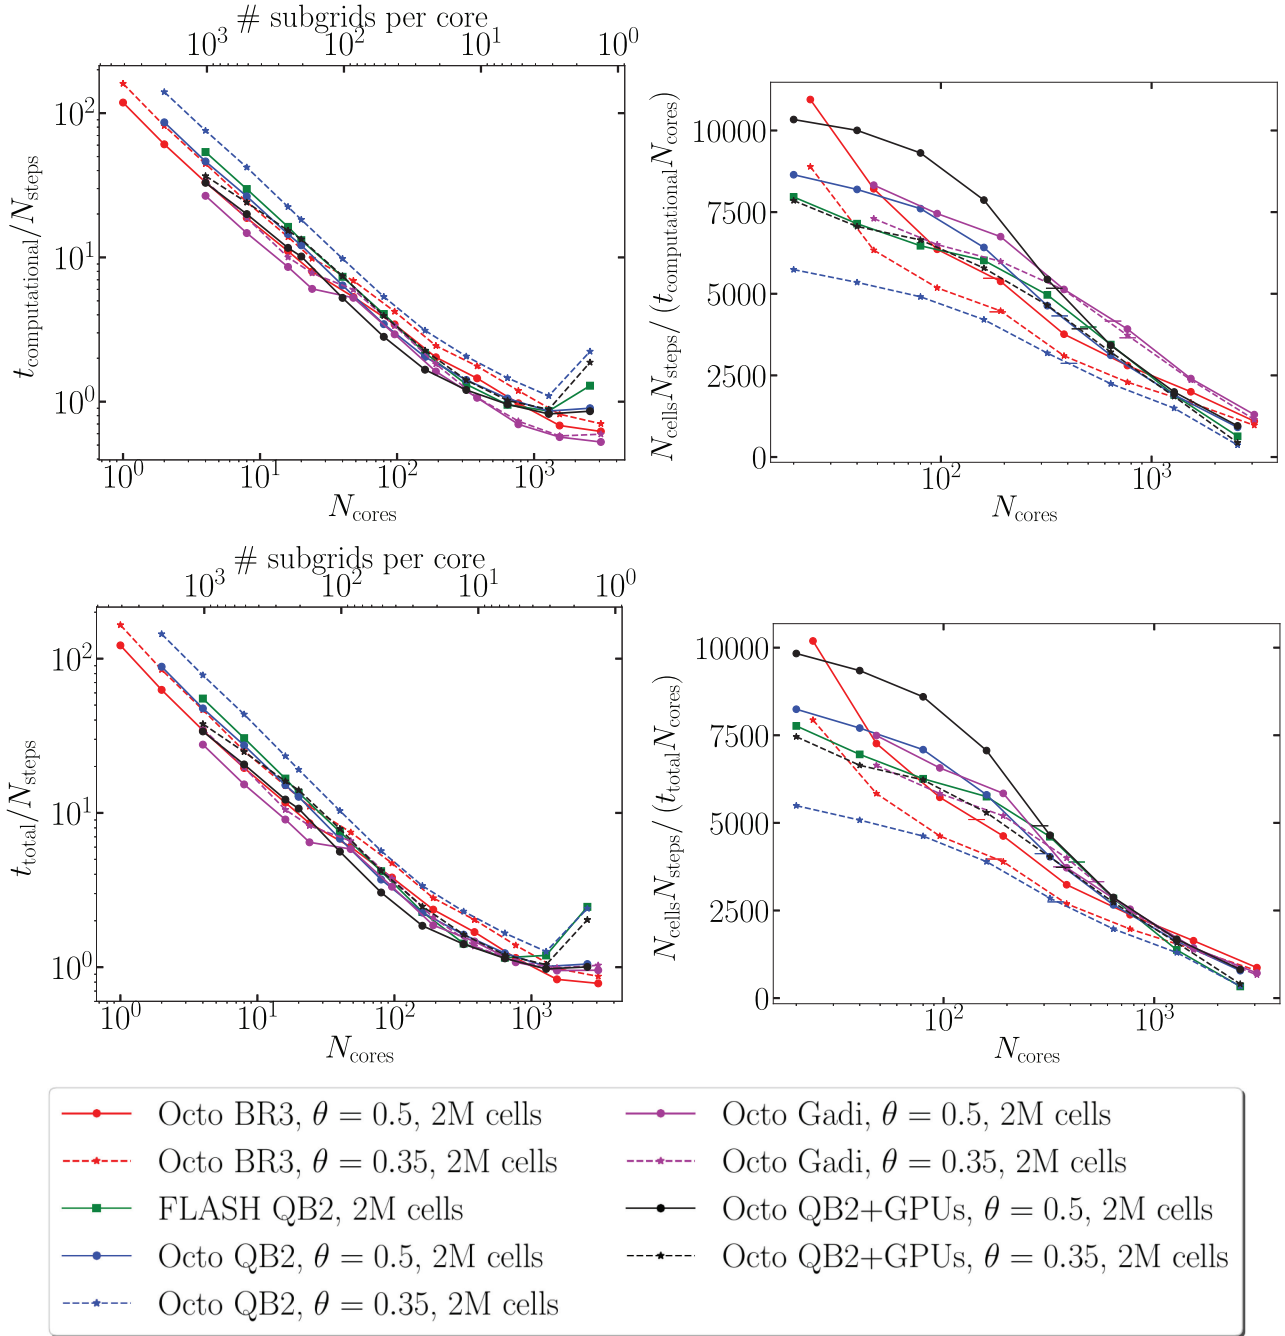


Figure 22. Scaling test carried out on the pulsating polytrope problem of a small size. QB2 refers to QueenBee2, BR3 refers to the BigRed3. The simulations contain $128^3 \approx 2M$ cells or $16^3 = 4096$ subgrids. The short horizontal segments mark the 0.5 efficiency.

in which case the communication between cores dominates the performance, and the scaling of the computational time becomes flat as well.

The values used for the figures are reported in Tables 7–9, organized by computer. In brackets, next to the values used for the plots, we calculate the ‘speedup’ ($S_n = (t_N/t'_N)/(N'_{\text{core}}/N_{\text{core}})$, where t_N and t'_N are the time-step lengths for runs that use N_{core} and N'_{core} cores, respectively). These efficiencies are calculated with respect to the run that uses an entire node in each machine (48 cores on Gadi, 20 cores on QB2 and 24 cores on BR3), when available or with the run with the least number of cores (N_{cores}), when not. On the RHS

panels of Figs 22 and 23, we mark the value of 0.5 efficiency by short horizontal segments over plotted on each curve.

In the tables, we have reported on the *total time* also known as wall-clock time, which includes the initialization, regridding, and the writing to disc part of the computation. If we had reported only the computational time, the efficiencies would increase.

We start by comparing the code performance for different values of the gravity solver parameter θ . This comparison was made with OCTO-TIGER with 2 million (solid lines) and 17 million cells (dashed lines) using BR3, QB2, QB2 + GPUS, and Gadi (red, blue, black, and magenta lines, respectively, Fig. 22). As expected, the smaller

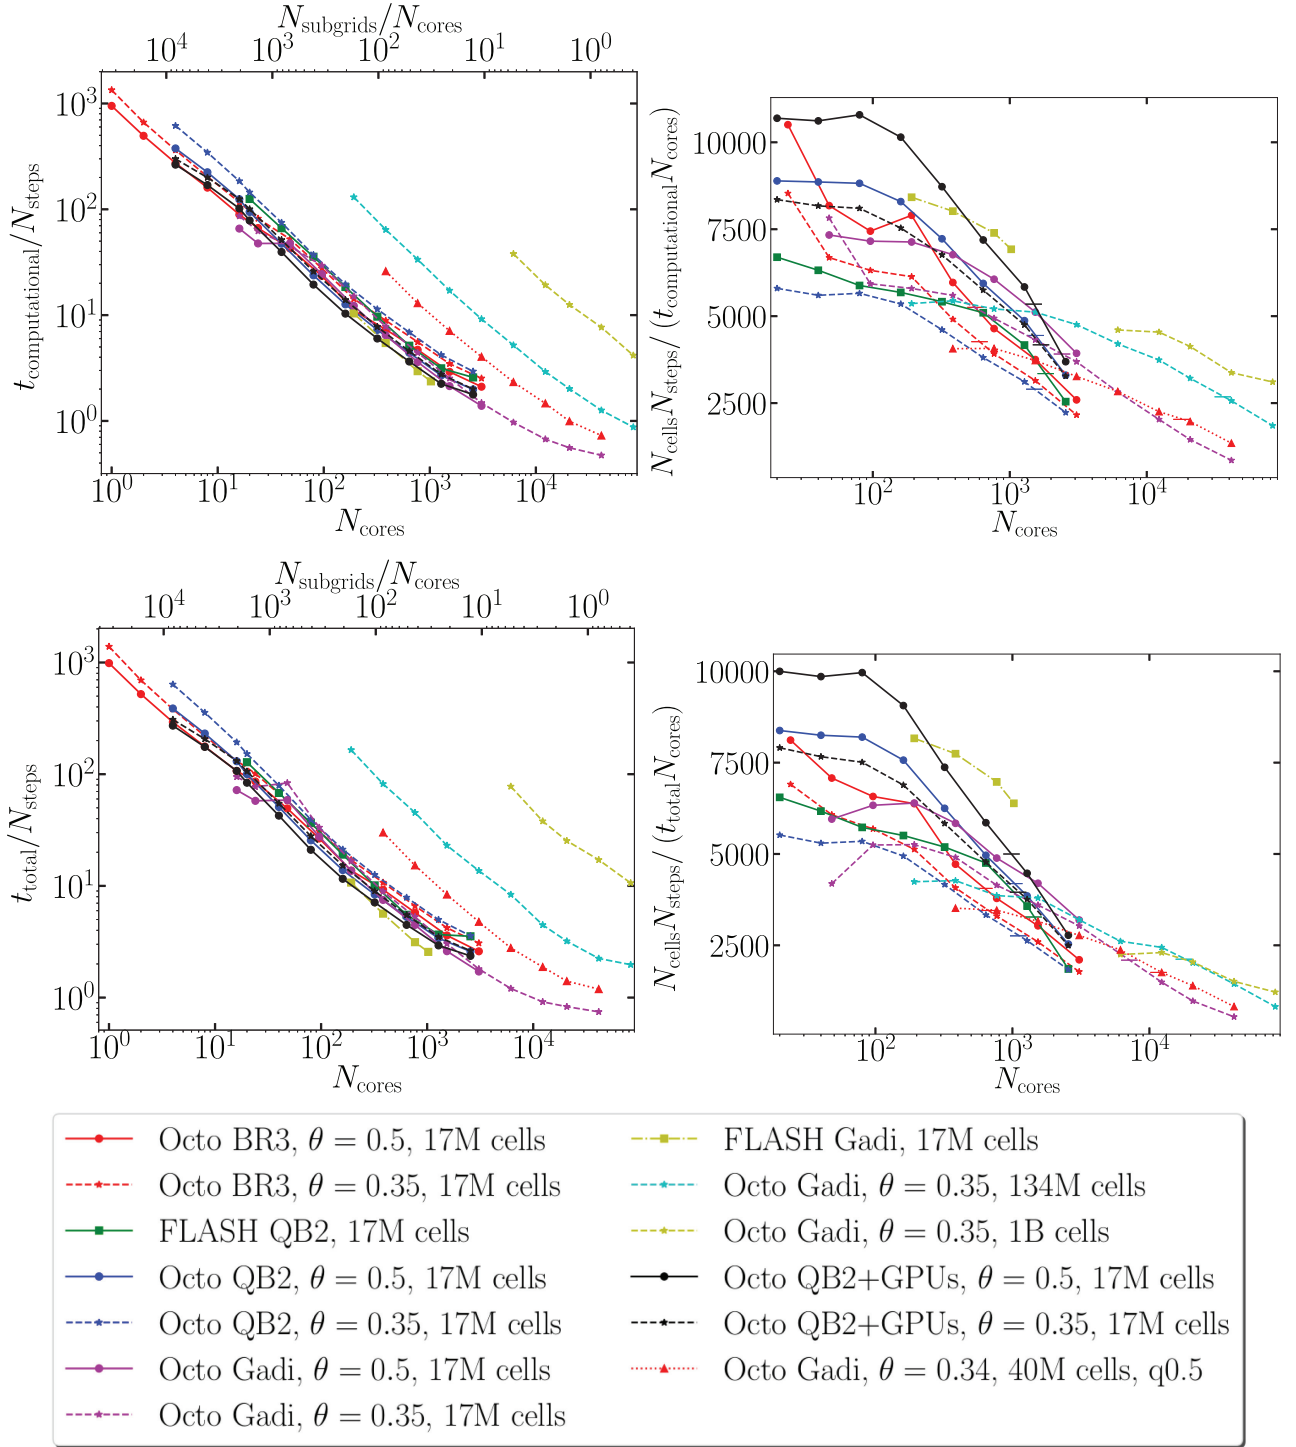


Figure 23. Scaling test carried out on the pulsating polytropic problem of big sizes. QB2 refers to QueenBee2, BR3 refers to the BigRed3. The simulations contain 256^3 (17M), 512^3 (134M), and 1024^3 (1B) cells or 4096, 32 768, 262 144 subgrids, respectively. One subgrids per core of the 17M cells simulations (top axis on left-hand panels) is 8 subgrids per core of the 134M cells, and 64 subgrids per core of the 1B cells simulations. The short horizontal segments mark the 0.5 efficiency.

the value of θ (the more precise the gravity solver solution) the longer the simulations take. At 2 million cells, the QB2 tests show an increase in the computational time by increasing the number of cores between 1280 and 2560, while this is less pronounced for the other two computers. This is due to a node-communication problem. The $\theta = 0.35$ simulations are doing more communication between

subgrids and hence nodes. The use of CPU + GPU improves the wall-clock time, but also displays the upturn. This behaviour is less pronounced for tests with more than 17 million cells (Fig. 23).

Next we compare the total time taken for the test for any code on any machine. For the 2 and 17-million cell simulations, the Gadi simulations are the fastest, although we note that for a range of core

Table 7. Wall-clock time per time-step in seconds for different scaling tests carried out on Gadi. In parenthesis, we report the speedup of a given calculation with respect to the calculation that uses one node (48 cores), interpreted as the fractional reduction of the time-step compared to the fractional increase in core count.

Code	OCTO-TIGER	OCTO-TIGER	FLASH	OCTO-TIGER	OCTO-TIGER	OCTO-TIGER	OCTO-TIGER	OCTO-TIGER ^d
N_{cells}	2M	2M	17M	17M	17M	134M	1000M	40M
θ	0.5	0.35	–	0.5	0.35	0.35	0.35	0.34
N_{cores}								
4	27.7	35.0	–	–	–	–	–	–
8	15.3	19.7	–	–	–	–	–	–
16	9.06	10.5	–	72.1	94.7	–	–	–
24	6.45	8.26	–	57.7	77.1	–	–	–
48	5.83 (–)	6.58 (–)	–	58.7 (–)	83.5 (–)	–	–	–
96	3.32 (88 per cent)	3.74 (88 per cent)	–	27.6 (125 per cent)	33.3 (125 per cent)	–	–	–
192	1.87 (78 per cent)	2.10 (78 per cent)	10.7 (–)	13.7 (126 per cent)	16.6 (125 per cent)	165 (–)	–	–
384	1.47 (50 per cent)	1.36 (60 per cent)	5.64 (95 per cent)	7.48 (117 per cent)	8.90 (117 per cent)	81.8 (101 per cent)	–	30.1 (–)
768	1.08 (34 per cent)	1.09 (38 per cent)	3.13 (85 per cent)	4.47 (99 per cent)	5.27 (99 per cent)	45.3 (91 per cent)	–	15.3 (98 per cent)
1024	–	–	2.57 (78 per cent)	–	–	–	–	–
1536	0.951 (19 per cent)	0.988 (21 per cent)	–	2.60 (86 per cent)	3.03 (86 per cent)	23.0 (90 per cent)	–	8.37 (90 per cent)
3072	0.955 (10 per cent)	1.03 (10 per cent)	–	1.71 (72 per cent)	1.80 (72 per cent)	13.6 (75 per cent)	–	4.77 (79 per cent)
6144	–	–	–	–	1.20 (55 per cent)	8.36 (62 per cent)	77.7 (–)	2.78 (68 per cent)
12 288	–	–	–	–	0.915 (36 per cent)	4.47 (58 per cent)	38.0 (102 per cent)	1.88 (50 per cent)
20 736	–	–	–	–	0.829 (23 per cent)	3.20 (48 per cent)	25.3 (91 per cent)	1.40 (40 per cent)
41 472	–	–	–	–	0.745 (13 per cent)	2.24 (34 per cent)	17.1 (67 per cent)	1.19 (23 per cent)
82 944	–	–	–	–	–	1.97 (19 per cent)	10.6 (54 per cent)	–

^dTest using a binary simulation with $q = 0.5$.

number between ~ 30 and ~ 300 it is the QB2 with GPU that is the fastest. This is particularly evident at 17 million cells, looking at the work per core per time plots (Fig. 23, RHS panels). Whether runs with GPUs are faster depends on whether the run is dominated by computation or by communication. GPU are computationally faster, but Gadi has a 4-time faster network communication than any other network used in this paper. Further information on OCTO-TIGER runs with GPUs can be found in Daiß et al. (2019b), and will be the subject of future papers.

We now come to the actual scaling properties. For the 2-million cell simulations, the efficiency of increasing the number of cores drops once we move past 1536 cores: on Gadi from 768 to 1546 to 3072 cores the efficiency drops from 34, to 19 to 10 per cent. This is similar on QB2 and BR3 (note that on QB2 we are comparing slightly different number of cores: 640, 1280, and 2560). At 17 million cells, the same three steps show efficiencies of 99, 86, and 72 per cent on Gadi, while for the other computers, they are markedly smaller, with, for instance, 47, 37, 26 per cent on BR3, and slightly better but similar values for QB2, making the runs on Gadi the fastest and most scalable.

Using GPUs with the 2- and 17-million cell simulations reduces the time per step between a factor of 1.5 and 2; the scalability of the hybrid CPU + GPU runs is slightly inferior to those of the pure CPU runs for the 2-million cell simulations, for which processors have less work to do, but almost identical to the non-GPU runs for the 17-million cell simulations.

Next, we look at scalability versus problem size by comparing runs with 2, 17, 134 million and 1 billion cells on Gadi. Looking at the efficiencies when increasing the number of cores between 1536 and 3072, we find 10, 72, and 75 per cent for 2, 17, and 134 million cells, respectively, which would imply that increasing the size of the problem from 17 to 134 million has not afforded us a much improved scaling. On the other hand, going from 6144 to 12 736 to 20 735 to 41 472 cores and finally to 82 944 cores, the efficiencies improve markedly upon doubling the number of cores for the 1 billion cell run compared to the 17-million cell run. In fact, it is remarkable that our 1 billion cell run is still scaling very well once we move to $\sim 80\,000$ cores, giving hope that we can carry out very high resolution simulations with reasonable wall-clock times (though the actual cost of the simulation in terms of CPU-hours would remain very high).

The scalability properties of FLASH were tested on QB2 at 2- and 17-million cell resolutions. They are very comparable to those of OCTO-TIGER. The computational time per step is intermediate between the $\theta = 0.35$ and 0.5, but tends to be longer than even the more accurate $\theta = 0.35$ OCTO-TIGER simulations when using a higher number of cores.

An additional scaling test was carried out using the first 30 time-steps of a binary interaction simulation similar to that of Section 4. The tests were to determine code scalability in view of AMR regridding. We performed a series of simulations on QB2, with increasing resolution by virtue of adding additional levels of refinement, see

Table 8. Wall-clock time per time-step in seconds for different scaling tests carried out on one stationary polytrope on QueenBee2. In parenthesis, we report the speedup of a given calculation with respect to the calculation that uses one node (20 cores), interpreted as the fractional reduction of the time-step compared to the fractional increase in core count.

Code	FLASH	OCTO	OCTO	OCTO	OCTO	FLASH	OCTO	OCTO	OCTO	OCTO
N_{cells}	2M	2M	2M	2M	2M	17M	17M	17M	17M	17M
θ	–	0.5	0.35	0.5	0.35	–	0.5	0.35	0.5	0.35
GPUs	–	–	–	✓	✓	–	–	–	✓	✓
N_{cores}										
2	–	88.7	144	–	–	–	–	–	–	–
4	55.1	47.6	78.1	33.7	37.7	–	388	637	272	308
8	30.5	27.4	43.6	20.6	24.9	–	231	356	175	207
16	16.7	15.1	23.4	12.2	16.0	–	131	193	107	132
20	13.5 (–)	12.7(–)	19.1 (–)	10.7 (–)	14.0 (–)	128 (–)	100 (–)	152 (–)	83.9 (–)	106 (–)
40	7.54 (90)	6.80 (93)	10.3 (93)	5.61 (95)	7.89 (89)	67.9 (94)	50.8 (98)	79.2 (96)	42.6 (99)	54.7 (97)
	per cent)	per cent)	per cent)	per cent)	per cent)	per cent)	per cent)	per cent)	per cent)	per cent)
80	4.19 (81)	3.70 (86)	5.67 (84)	3.05 (87)	4.21 (83)	36.6 (87)	25.6 (98)	39.2 (97)	21.0 (100)	27.9 (95)
	per cent)	per cent)	per cent)	per cent)	per cent)	per cent)	per cent)	per cent)	per cent)	per cent)
160	2.28 (74)	2.26 (70)	3.37 (71)	1.85 (72)	2.48 (71)	19.0 (84)	13.8 (90)	21.2 (90)	11.6 (91)	15.2 (87)
	per cent)	per cent)	per cent)	per cent)	per cent)	per cent)	per cent)	per cent)	per cent)	per cent)
320	1.42 (59)	1.62 (49)	2.30 (52)	1.41 (47)	1.63 (54)	10.1 (79)	8.39 (75)	12.6 (75)	7.11 (74)	8.98 (74)
	per cent)	per cent)	per cent)	per cent)	per cent)	per cent)	per cent)	per cent)	per cent)	per cent)
640	1.15 (37)	1.23 (32)	1.66 (36)	1.14 (29)	1.20 (37)	5.51 (73)	5.28 (59)	7.87 (60)	4.48 (59)	5.49 (60)
	per cent)	per cent)	per cent)	per cent)	per cent)	per cent)	per cent)	per cent)	per cent)	per cent)
1280	1.19 (18)	1.01 (20)	1.26 (24)	0.975 (17)	1.04 (21)	3.66 (55)	3.40 (46)	4.99 (48)	2.93 (45)	3.48 (48)
	per cent)	per cent)	per cent)	per cent)	per cent)	per cent)	per cent)	per cent)	per cent)	per cent)
2560	2.46 (4)	1.05 (9)	2.40 (6)	1.00 (8)	2.03 (5)	3.54 (28)	2.58 (30)	3.57 (33)	2.36 (28)	2.63 (32)
	per cent)	per cent)	per cent)	per cent)	per cent)	per cent)	per cent)	per cent)	per cent)	per cent)

Table 9. Wall-clock time per time-step in seconds for different scaling tests carried out on one stationary polytrope on BigRed3. In parenthesis, we report the speedup of a given calculation with respect to the calculation that uses one node (24 cores), interpreted as the fractional reduction of the time-step compared to the fractional increase in core count.

Code	OCTO-TIGER	OCTO-TIGER	OCTO-TIGER	OCTO-TIGER
N_{cells}	2M	2M	17M	17M
θ	0.5	0.35	0.5	0.35
N_{cores}				
1	122	165	987	1390
2	62.7	84.8	521	694
4	34.3	46.5	2938	382
8	19.5	25.2	177	219
16	11.7	15.0	107	132
24	8.57 (–)	11.0 (–)	86.11 (–)	101 (–)
48	6.01 (71 per cent)	7.48 (74 per cent)	49.4 (87 per cent)	57.5 (88 per cent)
96	3.81 (56 per cent)	4.72 (58 per cent)	26.6 (81 per cent)	30.7 (82 per cent)
192	2.36 (45 per cent)	2.81 (49 per cent)	13.8 (79 per cent)	17.0 (74 per cent)
384	1.69 (32 per cent)	2.03 (34 per cent)	9.25 (58 per cent)	10.7 (59 per cent)
768	1.15 (23 per cent)	1.38 (25 per cent)	5.77 (47 per cent)	6.61 (48 per cent)
1536	0.834 (16 per cent)	0.980 (18 per cent)	3.60 (37 per cent)	4.21 (38 per cent)
3072	0.785 (9 per cent)	0.872 (10 per cent)	2.59 (26 per cent)	3.07 (26 per cent)

Table 10 and Fig. 24. The code performance is aligned with that we have observed for the static polytrope on the same computer cluster (Table 8), although a side by side comparison is not possible in virtue because the number of cells was not exactly the same.

We finally carried out a last scaling test, where 12 regridding events took place in 100 time-steps, and where we started the simulation from an intermediate stage of the $q = 0.5$ simulation in order to trigger substantial regridding (Table 7, last column). This scaling runs were carried out on Gadi. Although the number of AMR boundaries varied by ≈ 30 per cent in these runs, the scaling was similar to a uniform grid with no regridding taking place (see red dotted triangle line in Fig. 23 and the last column in Table 7). We conclude that our regridding scheme scales well to high number of cores.

6 SUMMARY AND CONCLUSIONS

We have presented a thorough suite of benchmark tests and a comprehensive and detailed description of the AMR hydrodynamics code OCTO-TIGER. Although OCTO-TIGER has been used in tandem with other codes in previous publications in astrophysics (Kadam et al. 2016; Staff et al. 2018) as well as in computer science (e.g. Pfander et al. 2018; Daiß et al. 2019; Heller et al. 2019), there has never been an actual ‘method paper’.

OCTO-TIGER joins a small suite of codes specifically designed to study the merger of WD stars while also modelling the earlier, pre-merger phase. OCTO-TIGER is designed to be a scalable, accurate code that is able to exploit CPU and GPU computer architectures.

Table 10. Wall-clock time per time-step in seconds for different scaling tests carried out on QueenBee2 on a binary interaction simulation similar to the one of Section 4. In parenthesis, we report the speedup of a given calculation with respect to the calculation that uses one node (20 cores), except for the 45 million cell run, which is given with respect to the smallest number of cores tested (80 cores).

Code	OCTO-TIGER	OCTO-TIGER	OCTO-TIGER	OCTO-TIGER	OCTO-TIGER
N_{cells}	0.51M	1.5M	3.0M	10.5M	45M
θ	0.34	0.34	0.34	0.34	0.34
N_{cores}					
20	6.73	18.3	34.18	103	–
40	3.98 (85 per cent)	9.89 (92 per cent)	18.1 (95 per cent)	53.1 (97 per cent)	–
80	2.50 (67 per cent)	5.53 (83 per cent)	9.76 (88 per cent)	27.4 (94 per cent)	128
160	1.86 (45 per cent)	3.47 (66 per cent)	5.67 (76 per cent)	15.0 (86 per cent)	65.1 (98 per cent)
320	1.53 (27 per cent)	2.50 (46 per cent)	3.93 (54 per cent)	9.04 (71 per cent)	37.9 (84 per cent)
640	1.31 (16 per cent)	1.91 (30 per cent)	2.77 (30 per cent)	5.83 (55 per cent)	21.4 (75 per cent)
1280	1.19 (9 per cent)	1.56 (18 per cent)	2.11 (25 per cent)	4.02 (40 per cent)	12.6 (63 per cent)

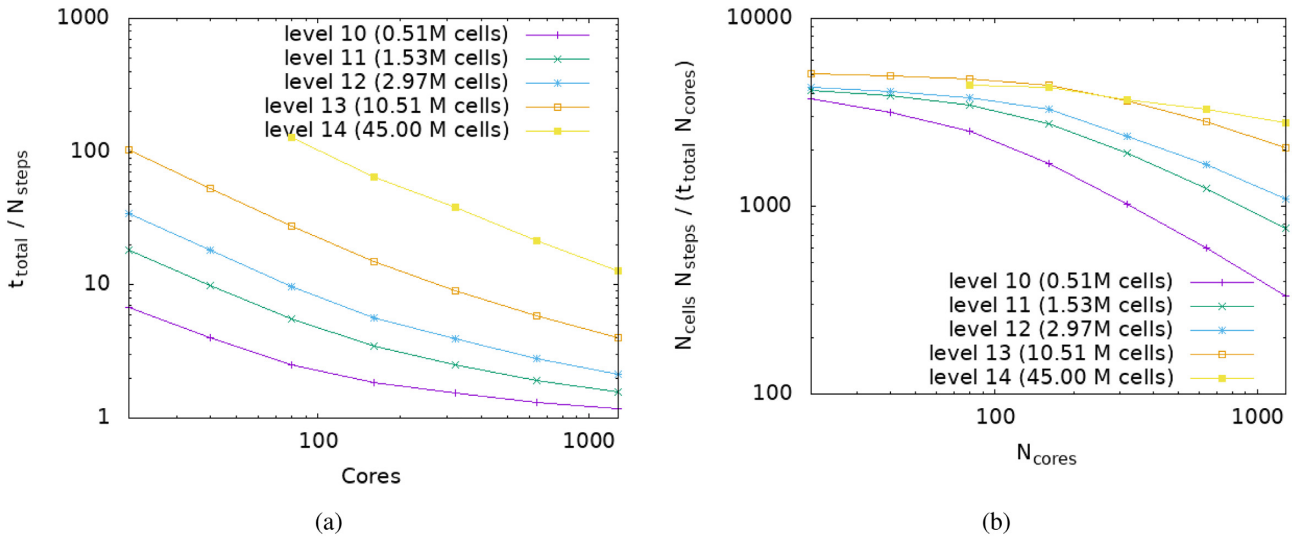


Figure 24. Scaling results for the first 30 time-steps of a binary simulation on LONI's QB2. The left-hand panel shows the total wall-clock time per time-step as a function of core count, while the right-hand panel shows the total throughput per core in terms of $N_{\text{cells}}N_{\text{time-steps}}/(t_{\text{wallclock}}N_{\text{cores}})$.

It is currently optimized to model stellar mergers for similar-sized stars that can be modelled using a star setup via a barotropic EoS, although for the evolution one can choose among a number of analytical EoS options. It does not currently include nuclear reactions such as for example the code CASTRO (Katz et al. 2016). OCTO-TIGER's gravity accuracy outperforms FLASH, with overall smaller residuals in the constant density sphere test (Section 3.4), although the comparison may not be with exactly the same parameters. The static and translating polytrope tests also show OCTO-TIGER's superior conservation of the peak density and overall mass of the structure, as well as the centre of mass position and overall velocity (for the translating polytropes; Sections 3.5 and 3.6).

In the rotating polytrope test, simulations in the rotating frame distinctly outperform those calculated in the inertial frame, although even in the inertial frame the polytrope retains its mass, position, and radius to excellent precision. OCTO-TIGER demonstrates a superior conservation of angular momentum and energy in this test, although the hydrodynamics is not well exercised in single grid tests where there is little gas motion (Section 3.7).

We have carried out scaling tests using a static polytrope, where the entire grid is fully refined with resolutions from 2 million to 1 billion cells. We have used three different computer clusters, two university-based mid-tier computers, QueenBee2 (Louisiana State University)

and BigRed3 (Indiana University) and one peak facility, Gadi, at the Australian Government's National Computational Infrastructure Centre. Strong scaling properties of the code are excellent up to $\sim 82\,000$, the largest number of cores tested, but only for the largest problems of 1 billion cells. With smaller problems of 134 million cells, performance improvement is obtained up to a maximum of $\sim 20\,000$ cores (after which the speedup is modest). For the 17 million cell test, similar to the number used in the binary simulation in Section 4, the optimum number of cores is ~ 6000 .

We have also carried out somewhat limited scaling tests of a binary problem using QueenBee2 and Gadi to measure the performance when regridding takes place and found that the scaling properties of the code are similar to those assessed using the stationary polytrope with no AMR regridding.

While the emphasis in this paper is not OCTO-TIGER's ability to use mixed CPU–GPU architectures, we have carried out a small test that showed that the use of 2 GPU units per node improves the step time by a factor of 1.5–2. Given the increased emphasis on the use of GPUs for this type of computation (see references in Section 2.10), we will further address this capability in future work.

Finally, we have used OCTO-TIGER to carry out intermediate resolution WD merger simulations with an ideal gas and a polytropic EoS and compared the results to a similar simulations carried out

by Motl et al. (2017) using two different codes, the grid code FLOWER, and the SPH code SNSPH (Fryer, Rockefeller & Warren 2006). The total run time of this 2–8 million cell simulation was approximately 8 d and it used 80 000 CPU-hours on 400 cores of the computer cluster QueenBee2. By comparison, the high-resolution cylindrical simulation in Motl et al. (2017), with an approximately similar resolution, ran for over a year. Because of the relatively small wall-clock times achieved by OCTO-TIGER thanks to its excellent scalability, higher resolution runs are possible.

Our improvements to OCTO-TIGER include (i) applying high resolution only to the binary through AMR (ii) exploiting the more favourable Courant condition from Cartesian coordinates while maintaining acceptable conservation, and (iii) leveraging a parallel framework that enables simulations on a very large number of computing elements. We have effectively reduced the wall-clock time for a merger simulation to complete from over a year to essentially one day. These developments enable a wide variety of numerical experiments that would not be possible or conclusive if run only at lower resolution.

An interesting question is how to best exploit resolution improvements. We know from Motl et al. (2017) that simply increasing the resolution will lower the initial mass transfer rate and greatly increase the early, pre-merger phase of the simulation. This would extend enormously the pre-merger time, making the simulation impossibly long. Under such circumstances, we would be forced to increase the artificial angular momentum extraction (the driving) to speed up the simulation to a time when the mass transfer rate is high enough that the merger takes place. As such, the mass transfer rate and the pre-merger time-scale cannot be measured with this type of simulation. On the other hand, there are other parameters of interest that would benefit from higher resolution. An example is the actual structure of the flow during the low-mass transfer rate time, or the details of the flow at or after the time of merger, including the determination of dredged-up elements.

In future papers, we will implement a number of additions. First, we will experiment with increasing the photospheric resolution, while at the same time implementing a radiation transport algorithm. With this the immediate aim is to model the light signature, something that would greatly benefit from sufficient resolution at the photosphere as thoroughly explained by Galaviz et al. (2017). After that, the following improvement will be to use model stars imported from 1D stellar structure simulations, instead of being constructed using barotropic solutions, and that can use a tabulated EoS. A follow-up study is already underway to simulate a $q = 0.1$ binary that mimics the V1309 Sco system – this would be only the second simulation of this system, after the 100 000 SPH particle hydrodynamic simulation of Nandez et al. (2014), and a $q = 0.7$ WD binary thought to be a typical progenitor system for the R Coronae Borealis stars (e.g. Clayton et al. 2011).

ACKNOWLEDGEMENTS

This work was supported by National Science Foundation Award 1814967. The numerical work was carried out using the computational resources (QueenBee2) of the Louisiana Optical Network Initiative (LONI) and Louisiana State University’s High Performance Computing (LSU HPC). Our use of BigRed3 at Indiana University was supported by Lilly Endowment, Inc., through its support for the Indiana University Pervasive Technology Institute. This research was undertaken with the assistance of resources (Gadi) from the National Computational Infrastructure (NCI Australia), an NCRIS enabled capability supported by the Australian Government. We wish to thank

Dale Roberts at NCI and Richard Miller at Macquarie University for helping the team to run scaling tests on an exorbitant number of Gadi cores. Finally, we thank Macquarie University for funding the Gadi computer time from the university allocation.

DATA AVAILABILITY

The source code of OCTO-TIGER (Marcello et al. 2021) is available on GitHub⁸ released under the Boost Software License Version 1. The build scripts to build OCTO-TIGER and its dependences are available on GitHub⁹ as well. The input files to run the OCTO-TIGER simulations are available on Zenodo (Marcello et al. 2020). Table 6 lists the software version used on the different clusters. Note that we used the pre-compiled MPI version on the cluster and therefore have some variation there. If the gcc compiler was recent enough to compile HPX and OCTO-TIGER, we used the pre-compiled version on the cluster.

For movies of the simulations and download of selected simulation data, see our YouTube channel.¹⁰

REFERENCES

- Amini P., Kaiser H., 2019, 2019 IEEE/ACM Third Annual Workshop on Emerging Parallel and Distributed Runtime Systems and Middleware (IPDRM), IEEE, Denver, CO. p. 26
- Bordner J., Norman M. L., 2012, Proceedings of the Extreme Scaling Workshop. BW-XSEDE ’12. University of Illinois at Urbana-Champaign, USA
- Bryan G. L., Norman M. L., Stone J. M., Cen R., Ostriker J. P., 1995, *Comput. Phys. Commun.*, 89, 149
- Bulla M., Sim S. A., Pakmor R., Kromer M., Taubenberger S., Röpke F. K., Hillebrandt W., Seitenzahl I. R., 2016, *MNRAS*, 455, 1060
- Burdge K. B., et al., 2020, *ApJ*, 905, L7
- Clayton G. C., 2012, *J. Am. Assoc. Var. Star Obs.*, 40, 539
- Clayton G. C., et al., 2011, *ApJ*, 743, 44
- Colella P., Woodward P. R., 1984, *J. Comput. Phys.*, 54, 174
- Copik M., Kaiser H., 2017, Proceedings of the 5th International Workshop on OpenCL, Association for Computing Machinery, New York NY. p. 1
- De Marco O., Long J., Jacoby G. H., Hillwig T., Kronberger M., Howell S. B., Reindl N., Margheim S., 2015, *MNRAS*, 448, 3587
- D’Souza M. C. R., Motl P. M., Tohline J. E., Frank J., 2006, *ApJ*, 643, 381
- Daiß G., et al., 2019, Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis. Association for Computing Machinery, Denver, Colorado, p. 62
- Daiß G. et al., 2019b, preprint (arXiv:1908.03121)
- Dehnen W., 2000, *ApJ*, 536, L39
- Després B., Labourasse E., 2015, *J. Comput. Phys.*, 290, 28
- Diehl P., Seshadri M., Heller T., Kaiser H., 2018, IEEE/ACM 4th International Workshop on Extreme Scale Programming Models and Middleware (ESPM2). Dallas, TX, USA, p. 19
- Duquennoy A., Mayor M., 1991, *A&A*, 248, 485
- Even W., Tohline J. E., 2009, *ApJS*, 184, 248
- Frank J., King A., Raine D. J., 2002, *Accretion Power in Astrophysics*, 3rd edn, Cambridge University Press, Cambridge, UK
- Fryer C. L., Rockefeller G., Warren M. S., 2006, *ApJ*, 643, 292
- Fryxell B. et al., 2000, *ApJS*, 131, 273
- Galaviz P., De Marco O., Passy J.-C., Staff J. E., Iaconi R., 2017, *ApJS*, 229, 36
- Hachisu I., 1986a, *ApJS*, 61, 479
- Hachisu I., 1986b, *ApJS*, 62, 461
- Heller T. et al., 2019, *Int. J. High Perform. Comput. Appl.*, 33, 699

⁸<https://github.com/STELLAR-GROUP/octotiger>

⁹<https://github.com/diehlpk/PowerTiger>

¹⁰<https://youtube.com/playlist?list=PL7vEgTL3FalaLHjQQH5UqEVR9-WkFeZp>

- Hillebrandt W., Niemeyer J. C., 2000, *ARA&A*, 38, 191
- Hurley M., Roberts P. H., Wright K., 1966, *ApJ*, 143, 535
- Ivanova N. et al., 2013, *A&AR*, 21, 59
- Ivezić Ž. et al., 2007, in Valsecchi G. B., Vokrouhlický D., Milani A., eds, Proc. IAU Symp. 236, Near Earth Objects, our Celestial Neighbors: Opportunity and Risk. Kluwer, Dordrecht, p. 353
- Jetley P., Gioachin F., Mendes C., Kale L. V., Quinn T., 2008, IEEE International Symposium on Parallel and Distributed Processing. p. 1
- Jha S. W., 2017, Type Iax Supernovae, Springer International Publishing AG, Basel, Switzerland. p. 375
- Kadam K., Motl P., Frank J., Clayton G., Marcello D. C., 2016, *MNRAS*, 462, 2237
- Kadam K., Motl P. M., Marcello D. C., Frank J., Clayton G. C., 2018, *MNRAS*, 481, 3683
- Kaiser H., Brodowicz M., Sterling T., 2009, International Conference on Parallel Processing Workshops, IEEE. p. 394
- Kaiser H., Heller T., Adelstein-Leibach B., Serio A., Fey D., 2014, Proceedings of the 8th International Conference on Partitioned Global Address Space Programming Models. ACM. p. 1
- Kaiser H. et al., 2020, *J. Open Source Softw.*, 5, 2352
- Kale L. V., Krishnan S., 1993, Technical Report, CHARM++: A Portable Concurrent Object Oriented System Based on C++, ACM. USA
- Kashyap R., Haque T., Lorén-Aguilar P., García-Berro E., Fisher R., 2018, *ApJ*, 869, 140
- Katz M. P., Zingale M., Calder A. C., Swesty F. D., Almgren A. S., Zhang W., 2016, *ApJ*, 819, 94
- Kippenhahn R., Kohl K., Weigert A., 1967, *Z. Astrophys.*, 66, 58
- Kippenhahn R., Thomas H.-C., Weigert A., 1968, *Z. Astrophys.*, 69, 265
- Kurganov A., Noelle S., Petrova G., 2000, *SIAM J. Sci. Comput.*, 23, 707
- MacLeod M., Loeb A., 2020, *ApJ*, 893, 106
- Marcello D. C., Tohline J. E., 2012, *ApJS*, 199, 35
- Marcello D., Kadam K., Clayton G. C., Frank J., Kaiser H., Motl P., 2016, Accretion Processes in Cosmic Sources, Proceedings of Science, Italy. p. 55
- Marcello D. C., Shiber S., Marco O. D., Frank J., Clayton G. C., Motl P. M., Diehl P., Kaiser H., 2020, *Files for Reproducing Results in Octo-Tiger: A New, 3D Hydrodynamic Code for Stellar Mergers That Uses HPX Parallelisation*, Zenodo, Switzerland, Available at: <https://zenodo.org/record/4432574#.YJQ5wyZOIH4>
- Marcello D. et al., 2021, GitHub, Available at: <https://github.com/STELLAR-GROUP/octotiger>
- Mason E., Diaz M., Williams R. E., Preston G., Bensby T., 2010, *A&A*, 516, A108
- Meyer F., Meyer-Hofmeister E., 1979, *A&A*, 78, 167
- Motl P. M., Tohline J. E., Frank J., 2002, *ApJS*, 138, 121
- Motl P. M., Frank J., Tohline J. E., D'Souza M. C. R., 2007, *ApJ*, 670, 1314
- Motl P. M., Frank J., Staff J., Clayton G. C., Fryer C. L., Even W., Diehl S., Tohline J. E., 2017, *ApJS*, 229, 27
- Nandez J. L. A., Ivanova N., Lombardi J. C., Jr, 2014, *ApJ*, 786, 39
- Pakmor R., Kromer M., Taubenberger S., Sim S. A., Röpke F. K., Hillebrandt W., 2012, *ApJ*, 747, L10
- Pejcha O., Metzger B. D., Tomida K., 2016, *MNRAS*, 461, 2527
- Pfander D., Daiß G., Marcello D., Kaiser H., Pflüger D., 2018, Proceedings of the International Workshop on OpenCL. IWOCCL '18. ACM, New York, p. 19
- Reichardt T. A., De Marco O., Iaconi R., Tout C. A., Price D. J., 2019, *MNRAS*, 484, 631
- Ricker P. M., Timmes F. X., Taam R. E., Webbink R. F., 2019, in Oskinova L. M., Bozzo E., Bulik T., Douglas R.G., eds, Proc. IAU Symp. 346, IAUS 346: High-mass X-ray binaries: illuminating the passage from massive binaries to merging compact objec, Cambridge University Press, Cambridge UK. p. 449
- Rucinski S., 2010, in Kalogera V., van der Sluys M., eds, AIP Conf. Proc., Vol. 1314, INTERNATIONAL CONFERENCE ON BINARIES: In celebration of Ron Webbink's 65th Birthday, AIP Publishing, Melville, NY, p. 29
- Schneider F. R. N., Ohlmann S. T., Podsiadlowski P., Röpke F. K., Balbus S. A., Pakmor R., Springel V., 2019, *Nature*, 574, 211
- Sedov L. I., 1946, *J. Appl. Math. Mech.*, 10, 241
- Shu C.-W., Osher S., 1989, *J. Comput. Phys.*, 83, 32
- Smith R. C., 1984, *QJRAS*, 25, 405
- Sod G. A., 1978, *J. Comput. Phys.*, 27, 1
- Solheim J.-E., 2010, *PASP*, 122, 1133
- Staff J. E. et al., 2018, *ApJ*, 862, 74
- Tauris T. M., van den Heuvel E. P. J., 2014, *ApJ*, 781, L13
- The C++ Standards Committee, 2017, *Technical Report*, ISO International Standard ISO/IEC 14882:2017, Programming Language C++. International Organization for Standardization (ISO), Geneva
- The C++ Standards Committee, 2020, Technical Report, ISO International Standard ISO/IEC 14882:2020, Programming Language C++. International Organization for Standardization (ISO), Geneva
- Thoman P. et al., 2018, *J. Supercomput.*, 74, 1422
- Tylenda R. et al., 2011, *A&A*, 528, A114
- Verbunt F., Rappaport S., 1988, *ApJ*, 332, 193
- Warner B., 2003, Cataclysmic Variable Stars, Cambridge University Press, Cambridge, UK
- Webbink R. F., 1992, in van den Heuvel E. P. J., Rappaport S. A., eds, NATO Advanced Study Institute (ASI) Series C Vol. 377, X-Ray Binaries and Recycled Pulsars, Cambridge University Press, Cambridge, UK. p. 269

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