

RESEARCH

Open Access



Entropy-limited hydrodynamics: a novel approach to relativistic hydrodynamics

Federico Guercilena^{1*}, David Radice^{2,3} and Luciano Rezzolla^{1,4}

Abstract

We present entropy-limited hydrodynamics (ELH): a new approach for the computation of numerical fluxes arising in the discretization of hyperbolic equations in conservation form. ELH is based on the hybridisation of an unfiltered high-order scheme with the first-order Lax-Friedrichs method. The activation of the low-order part of the scheme is driven by a measure of the locally generated entropy inspired by the artificial-viscosity method proposed by Guermond et al. (*J. Comput. Phys.* 230(11):4248-4267, 2011, doi:10.1016/j.jcp.2010.11.043). Here, we present ELH in the context of high-order finite-differencing methods and of the equations of general-relativistic hydrodynamics. We study the performance of ELH in a series of classical astrophysical tests in general relativity involving isolated, rotating and nonrotating neutron stars, and including a case of gravitational collapse to black hole. We present a detailed comparison of ELH with the fifth-order monotonicity preserving method MP5 (Suresh and Huynh in *J. Comput. Phys.* 136(1):83-99, 1997, doi:10.1006/jcph.1997.5745), one of the most common high-order schemes currently employed in numerical-relativity simulations. We find that ELH achieves comparable and, in many of the cases studied here, better accuracy than more traditional methods at a fraction of the computational cost (up to ~50% speedup). Given its accuracy and its simplicity of implementation, ELH is a promising framework for the development of new special- and general-relativistic hydrodynamics codes well adapted for massively parallel supercomputers.

Keywords: flux limiters; entropy limited hydrodynamics; numerical methods; central schemes

1 Introduction

Large-scale general-relativistic hydrodynamical numerical simulations have been shown to be a very powerful tool for the study of astrophysical systems involving compact objects such as black holes and neutron stars (Font 2008; Shibata and Taniguchi 2011; Rezzolla and Zanotti 2013; Martí and Müller 2015; Shibata 2016; Baiotti and Rezzolla 2016; Paschalidis 2016). The realisation of such simulations requires dealing with very different physical, mathematical and computational issues. One of the most challenging of such issues, and one that could lead to significant differences on the outcome of resolution-limited simulations, is the choice of the numerical method for the solution of the hydrodynamics equations (one such difference would be e.g., the dephasing of the gravitational waveforms in

binary neutron star merger simulations, where different numerical schemes can lead to different dynamics of the matter bulk, see e.g., Baiotti et al. 2011; Read et al. 2013; Radice et al. 2014a; Thierfelder et al. 2011; Bernuzzi et al. 2012).

The most commonly used methods in this context are collectively known as high-resolution shock-capturing (HRSC) techniques (see Rezzolla and Zanotti 2013 for an introduction). Belonging to this class, which contains both finite-differences and finite-volumes schemes, are e.g., slope limiting methods (e.g., Roe 1985), the piece-wise parabolic method (PPM) (Colella and Woodward 1984; Martí and Müller 1996), the fifth-order monotonicity preserving (MP5) method (Suresh and Huynh 1997), essentially/weighted non-oscillatory (ENO/WENO) methods (Harten et al. 1987; Liu et al. 1994; Jiang and Shu 1996; De Pietri et al. 2016; Bernuzzi and Dietrich 2016) and many others.

HRSC methods are very effective in dealing with shocks and suppressing spurious oscillations, and have been em-

* Correspondence: guercilena@th.physik.uni-frankfurt.de

¹Institut für Theoretische Physik, Goethe Universität, Max-von-Laue-Str. 1, Frankfurt am Main, 60438, Germany

Full list of author information is available at the end of the article

ployed with varying degree of success in astrophysical simulations. In recent times much work has gone into improving these schemes (e.g., by innovative mesh refinement techniques such as in DeBuhr et al. 2015) or moving beyond them; one promising alternative paradigm is that of discontinuous Galerkin methods (see, e.g., Radice and Rezzolla 2011; Bugner et al. 2016; Zanotti et al. 2015; Kidder et al. 2016; Miller and Schnetter 2016). Many of such schemes, however, potentially suffer from a few shortcomings: (i) they are complex to derive and implement, or to extend and modify (e.g., to increase the formal order of accuracy); (ii) they often depend on many coefficients that require some degree of optimisation (e.g., the WENO methods); (iii) they can lead to load imbalance in parallel implementations as a result of their complexity.

In this work we propose a different approach that is able to address some of these shortcomings (especially the points (i) and (iii) previously mentioned), which we refer to as ‘entropy-limited hydrodynamics’ (ELH) and formulate it in a finite-differences framework. The underlying concept is relatively straightforward: the hydrodynamical fluxes are computed using an unfiltered, high-order stencil, to which a contribution from a low-order, stable numerical flux (the Lax-Friedrichs flux) is added in order to ensure stability. To determine which gridpoints are in need of the low-order contribution, we employ a ‘shock detector’, which not only marks region of the computational domain requiring the limiter, but also determines the relative weights of the high and low-order fluxes.

The use of a hybrid numerical flux to achieve both accuracy and stability places our method in the class of flux-limiting schemes (see, e.g., the classification in Leveque 1992), which have long been a feature in the panorama of numerical schemes for hydrodynamics. In this context the Lax-Friedrichs flux is a common choice for low-order methods, being monotone and dissipative (a different example of combining high- and low-order methods, in the context of the reconstruction method, would be e.g., Tchekhovskoy et al. 2007).

To drive the activation of the Lax-Friedrichs method, a criterion to flag generically problematic points of the computational domain is needed. Such a criterion is offered by the entropy viscosity function described by Guermond et al. (we refer primarily to Guermond et al. 2011, but see also Guermond and Pasquetti 2008; Zingan et al. 2013), in which the local production of entropy is used to identify shocks. Since entropy is produced only in the presence of shocks, this results in a stable method able to recover high-order in regions of smooth flow. We extend the definition of the entropy viscosity from the classical to the relativistic case, and employ it to drive the flux limiting scheme rather than as a weight to additional viscous terms in the hydrodynamical equations. As a result, and in contrast to the approach by Guermond et al. (2011), we do not modify

the underlying equations of relativistic hydrodynamics by introducing additional entropy-related terms.

In the following we describe the method and the details of our implementation, then report the results of tests we conducted in order to gauge its behaviour against a standard HRSC method, namely, MP5 (Suresh and Huynh 1997). The paper is structured as follows: in Section 2 we briefly summarise the equations of relativistic hydrodynamics and the finite-differences framework we employ. In Section 3 the ELH method and its implementation are described, while the results of the numerical tests are collected in Section 4. We present our conclusions and outlook in Section 5.

In the following we use the spacetime signature $(-, +, +, +)$, with Greek indices running from 0 to 3 and Latin indices from 1 to 3. We also employ the Einstein convention for the summation over repeated indices. Unless otherwise stated, all quantities are expressed in a geometrized system of units in which $c = G = 1$.

2 Relativistic hydrodynamics: theory and numerics overview

2.1 Relativistic hydrodynamics

We summarize here the mathematical framework of relativistic hydrodynamics. In the interest of simplicity, the discussion is limited to special relativity, while the general-relativistic case, which is relevant for the neutron-star tests of Section 4.2, can be found in Appendix 1.

Since most of our interest is in modelling neutron-star matter, we assume it to be described by a perfect fluid, therefore with a corresponding energy-momentum tensor given by

$$T_{\mu\nu} = \rho h u_{\mu} u_{\nu} + p \eta_{\mu\nu}, \quad (1)$$

where $\eta_{\mu\nu}$ is the Minkowski metric, ρ is the rest-mass density, u^{μ} is the fluid four-velocity, p is the pressure, $h = 1 + \epsilon + p/\rho$ is the specific enthalpy and ϵ the specific internal energy (Rezzolla and Zanotti 2013). The equations of motion for the fluid are the conservation of the stress-energy tensor

$$\partial_{\mu} T^{\mu\nu} = 0, \quad (2)$$

and conservation of rest mass

$$\partial_{\mu} (\rho u^{\mu}) = 0. \quad (3)$$

These two sets of equations are closed by an equation of state (EOS) $p = p(\rho, \epsilon)$, and we here assume a simple ideal-fluid (or Gamma-law) EOS:

$$p = \rho \epsilon (\Gamma - 1), \quad (4)$$

with Γ the adiabatic index.

Equations (2) and (3) can be cast in conservation form and therefore written symbolically as (Martí and Müller 2015)

$$\partial_t \mathbf{U} + \partial_i \mathbf{F}^i = \mathbf{S}, \quad (5)$$

by defining the conserved variables \mathbf{U} as

$$\begin{aligned} \mathbf{U} &:= \begin{pmatrix} D \\ S_j \\ \tau \end{pmatrix} \\ &:= \begin{pmatrix} \rho W \\ \rho h W^2 v_j \\ \rho h W^2 - p - \rho W \end{pmatrix}. \end{aligned} \quad (6)$$

These are functions of the ‘primitive’ variables (ρ, v_i, p, ϵ) . In these expressions and in the following the fluid three-velocity measured by the normal (or Eulerian) observers is defined as $v^i := u^i/W$ and the Lorentz factor is $W := (1 - v^i v_i)^{-1/2} = u^t$. The fluxes \mathbf{F}^i are

$$\mathbf{F}^i = \begin{pmatrix} v^i D \\ S_j v^j + p \delta_j^i \\ S^i - D v^i \end{pmatrix}, \quad (7)$$

where δ_j^i is the Kronecker delta.

Note that the source functions \mathbf{S} are identically zero in special relativity, but this is no longer the case in a generic spacetime, where metric-dependent terms appear both in the fluxes and sources (see Appendix 1).

2.2 Numerical methods

The ELH method proposed here has been implemented in the code `WhiskyTHCEL` as a variant of the `WhiskyTHC` code (Radice and Rezzolla 2012; Radice et al. 2014a, 2014b) based on the `Einstein Toolkit` (Löffler et al. 2012; Zilhão and Löffler 2013; Einstein Toolkit 2010). `WhiskyTHC` implements both finite-difference and finite-volume methods applied to a characteristic-variables decomposition with a Lax-Friedrichs flux-splitting for upwinding. It also crucially provides a positivity preserving limiter to cope with large rest-mass density jumps, e.g., as those appearing across the surface of compact stars. In the following we summarise the main components of the underlying algorithm and refer the interested reader to Radice and Rezzolla (2012) and Radice et al. (2014a) for a more detailed description.

Given a discrete mesh with Δ being the spatial grid spacing, the finite-difference algorithm we employ provides an estimate for the right-hand-side of an evolution equation in flux-conservative form as

$$\partial_t U|_i = -\frac{f_{i+1/2} - f_{i-1/2}}{\Delta} + S_i, \quad (8)$$

i.e., as a difference between numerical fluxes at the cell interfaces, plus the sources contribution (to simplify the notation, here U is any one of the components of (6)).

The numerical fluxes $f_{i\pm 1/2}$ are obtained via a reconstruction operator, i.e., an operator yielding a high-order approximation of a generic function at a given point using its volume averages (see, e.g., Leveque 1992; Rezzolla and Zanotti 2013; Martí and Müller 2015). Out of the variety of reconstruction operators available in the literature and implemented in `WhiskyTHC`, we focus here on the fifth-order and seventh-order unfiltered stencils

$$\begin{aligned} {}^5S^- &:= \frac{1}{60}(2f_{i-2} - 13f_{i-1} + 47f_i \\ &\quad + 27f_{i+1} - 3f_{i+2}) \end{aligned} \quad (9)$$

and

$$\begin{aligned} {}^7S^- &:= -\frac{1}{140}f_{i-3} + \frac{5}{84}f_{i-2} - \frac{101}{420}f_{i-1} + \frac{319}{420}f_i \\ &\quad + \frac{107}{210}f_{i+1} - \frac{19}{210}f_{i+2} + \frac{1}{105}f_{i+3}, \end{aligned} \quad (10)$$

returning the value of the flux at $x_{i+1/2}$, and which we refer to as U5 and U7, respectively (we have only written here the left-biased operators S^- , since the right-biased ones S^+ are symmetric).^a

Furthermore, we select the MP5 scheme as our benchmark against which to test the properties of the EL method. MP5 is built on top of the U5 stencil, but the resulting fluxes are limited so as to preserve monotonicity near discontinuities (Suresh and Huynh 1997; Mignone et al. 2010; Radice and Rezzolla 2012). MP5 offers a good compromise between robustness and accuracy and it has been successfully employed in several realistic scenarios in which it also achieved high convergence-order (Radice et al. 2014a, 2014b; Radice et al. 2015). It has therefore become a *de facto* standard in our work, hence we use it as a reference.

To ensure the stability of the scheme, the reconstruction must be appropriately upwinded, i.e., a right- (left-) biased operator has to be applied to the left- (right-) going part of the flux. We therefore split the flux f in a right-going flux f^+ and a left-going one f^- , so that $f = f^+ + f^-$. The splitting is performed using the Lax-Friedrichs or Rusanov flux splitter (Shu 1997), i.e.,

$$f^\pm := f(U) \pm \kappa U, \quad \kappa := \max |f'(U)|, \quad (11)$$

where the maximum is taken over the stencil of the reconstruction operator.

The reconstruction procedure outlined above can be applied on each equation in the system (5) (this is also called a components split) or to its local characteristic variables (in which case it is referred to as a characteristics split).

A further ingredient in our algorithm is a so-called ‘positivity-preserving’ limiter (Hu et al. 2013; Radice et al. 2014a). The basic idea is to split the numerical flux in Eq. (8) in two contributions, as:

$$f_{i+1/2} = \theta f_{i+1/2}^{\text{HO}} + (1 - \theta) f_{i+1/2}^{\text{LF}}, \quad (12)$$

where f^{HO} is the original high-order flux, while the Lax-Friedrichs flux f^{LF} has the standard form

$$f_{i+1/2}^{\text{LF}} := \frac{1}{2}(f_i + f_{i+1}) - \frac{\kappa}{2}(U_i - U_{i+1}) \quad (13)$$

and κ is defined as in Eq. (11). For a single Euler step, the result of the evolution of U can be explicitly written as

$$U_i^{n+1} = \frac{1}{2}(U_i^n + 2\lambda f_{i-1/2}) + \frac{1}{2}(U_i^n - 2\lambda f_{i+1/2}), \quad (14)$$

where the fluxes are defined as in Eq. (12) and λ depends on the maximum propagation speed of the system as well as the CFL factor. The value of θ is defined as the one that makes both terms of Eq. (14) positive. Applied to the continuity equation this guarantees that the density never becomes negative (see Wu 2017 for a way to generally ensure the physicality of the fluid state-vector in a generic spacetime).^b

We note that this algorithm does not free us from having to employ an artificial floor (or atmosphere) to treat (ideally) vacuum regions: these are filled with a uniform and tenuous fluid with rest-mass density ρ_{atmo} . Whenever in the subsequent evolution the rest-mass density of a grid-point falls below the floor value ρ_{atmo} , it is reset to the floor value, its three-velocity is set to zero and the specific internal energy is set to the corresponding value coming from the EOS. In neutron star simulations we fix the floor at $\rho_{\text{atmo}} = 10^{-16} M_{\odot}^{-2}$, i.e. the typical value of $\rho_{\text{atmo}}/\rho_{\text{max}}$ is $\sim 10^{-13}$.

Details of the algorithms we employ to evolve the spacetime and couple it to the fluid evolution are given in Appendix 2.

The last step in the algorithm is the actual time evolution. Since after the spatial discretization, the original PDEs to be solved are in the form of a coupled systems of ODEs, this is taken care of in a method-of-lines (MOL) fashion by means of a fixed step Runge-Kutta time integrator. We employ either the standard fourth-order Runge-Kutta method (RK4) or a third-order (RK3, see Gottlieb et al. 2009) method with strong stability preserving (SSP) properties.

3 Entropy-limited hydrodynamics

3.1 Description of the scheme

The scheme we propose consists of two building blocks: (1) a function detecting shocks; (2) a limiter scheme of the high-order fluxes. We will start discussing the latter.

As customary in flux-limiting schemes, we modify the high-order approximation of the flux by combining (or ‘hybridising’) it with a (local) Lax-Friedrichs flux contribution as in (12).

Of course, the hybridisation of the high-order flux with the Lax-Friedrichs should be activated only in regions of the flow that are problematic. In order to flag such regions we introduce a regularisation function that we refer to as the ‘viscosity’ ν . Hence, we *redefine* the parameter $\theta \in [0, 1]$ in Eq. (12) in terms of the quantity

$$\theta := \min \left[\tilde{\theta}, 1 - \frac{1}{2}(\nu_i + \nu_{i+1}) \right], \quad (15)$$

so that the contribution of the Lax-Friedrichs flux grows linearly with the viscosity. The value of the coefficient $\tilde{\theta}$ is the one mentioned in the previous section to guarantee the positivity of the rest-mass density. With the choice (15) for the limiting coefficient θ , additional dissipation is inserted when ν attains large values as well as in near-vacuum regions. On the other hand, in regions where the flow is smooth and away from near vacuum, θ is close to unity, ensuring the use of the high-order flux and preserving the high accuracy of the method.

We still need to associate the viscosity ν to some property of the flow. To this scope we take inspiration from the work of Guermont and collaborators (Guermont et al. 2011) and associate ν to the specific entropy s . In general, the precise functional form of s will depend on the EOS, but for the simpler case of a perfect fluid with an ideal-fluid EOS (cf., Eq. (4)), the specific entropy can be shown to be equal to (apart from constant multiplicative factors) to (Rezzolla and Zanotti 2013)

$$s = \log \left(\frac{\epsilon}{\rho^{\Gamma-1}} \right). \quad (16)$$

Of course, the specific entropy must satisfy the second law of thermodynamics, so that we can introduce the *entropy residual*, or entropy-production rate, \mathcal{R} as

$$\mathcal{R} := \partial_{\mu}(\rho s u^{\mu}) \geq 0. \quad (17)$$

As a result, the computation of the entropy residual \mathcal{R} , effectively represents the first step in defining the viscosity and hence the root to limiter parameter θ .

The expected behaviour of the entropy residual is that it cannot decrease in time and that is spatially restricted to very small regions in the neighbourhood of shocks, ideally expressed a delta function peaked at the location \mathbf{x}_s of shocks, i.e., $\mathcal{R} \propto \delta(\mathbf{x} - \mathbf{x}_s)$. A physical justification for this latter expectation is rather simple to motivate. Euler equations generally apply to perfect fluids, and while they can capture non-ideal features (i.e., shocks), the description of the latter is only approximate. As long as the flow is

smooth and the perfect-fluid approximation holds, all phenomena are reversible and there can be no production of entropy. However, in those regimes where the perfect-fluid approximation breaks down and non-ideal effects appear, namely, at the location of shocks, the entropy production is nonzero and the entropy jumps locally to a higher value. Since shocks are regions of dimension $N - 1$ in spatial manifolds with N spatial dimensions, the entropy residual \mathcal{R} must be a Dirac delta peaked at shock locations for it to provide a finite contribution.

To seal the strict connection between the entropy viscosity ν_e and the entropy residual and to embody the property that this quantity should be a function of the spatial discretization, we define it as

$$\nu_e := c_e \Delta |\mathcal{R}|. \quad (18)$$

The absolute value is used for the entropy residual since the inequality (18) is not guaranteed to be satisfied during the numerical integration. In fact, \mathcal{R} , having to approximate a delta function, is expected to show an oscillatory behaviour with potential negative values in practical numerical applications. In expression (18), Δ is the spacing of the mesh, c_e is a positive tunable constant with dimensions of $[\text{time}]^3 \times [\text{temperature}] \times [\text{mass}]^{-1}$, so that ν_e is effectively dimensionless. We note that despite ν_e not having the dimensions of a viscosity, we still refer to as the ‘entropy viscosity’, mostly for convenience and in analogy with the very similar quantity defined in Guermond et al. (2011).

An additional benefit of this definition is the ability of the resulting scheme in differentiating automatically between shocks and contact discontinuities. This follows from the fact that at contact discontinuities there is no entropy production and therefore the viscosity there would be zero as well (Rezzolla and Zanotti 2013).

A potential problem of the definition (18) is that it can lead to an unbounded value since the entropy residual \mathcal{R} is not physically upper limited. In our case the value of θ cannot however exceed unity, and so the viscosity must not exceed this value as well. To enforce this requirement and cut-off excessively large values of the entropy viscosity we set the entropy viscosity to be used in the limiter (15) as

$$\nu := \min[\nu_e, c_{\max}], \quad (19)$$

where ν_e is given by Eq. (18). Here, c_{\max} is a tunable dimensionless coefficient, which, together with the other tunable coefficient c_e , we have assumed to be equal to one in all of the tests presented here. As we will comment in Section 4, the results are not very sensitive to the choices made for these coefficients.

3.2 Numerical implementation

In our numerical implementation we compute the entropy residual (17) by first rewriting its definition in a way that involves only derivatives of the specific entropy s

$$\mathcal{R} = \partial_\mu (s \rho u^\mu) = s \partial_\mu (\rho u^\mu) + \rho u^\mu \partial_\mu s = \rho u^\mu \partial_\mu s, \quad (20)$$

where the continuity equation (3) was used to obtain the final expression in (20). By expressing the 4-velocity u^μ in terms of the fluid three-velocity v^i , we finally write the residual as

$$\mathcal{R} = \rho W (\partial_t s + v^i \partial_i s). \quad (21)$$

The spatial derivatives in Eq. (21) are approximated with a standard centered finite-difference stencil of order $p + 1$, where p is the order of the stencil used to approximate the physical fluxes. This restriction arises from the need to ensure that the viscosity converges to zero fast enough not to spoil the overall convergence of the scheme at the nominal order. The time derivative in (21) is also approximated by finite differencing. In particular, at every iteration we use the current value of the specific entropy and the values at the two previous timesteps to compute a second-order approximation of $\partial_t s$ via a one-sided stencil, i.e., as

$$(\partial_t s)^n = \frac{1}{2\Delta t} (3s^n - 4s^{n-1} + s^{n-2}) + \mathcal{O}((\Delta t)^2). \quad (22)$$

A few remarks are useful at this point. First, the time derivative of the specific entropy in Eq. (21) is computed with a low-order method and this could in principle be a limiting factor for the convergence properties of the overall scheme. In practice, however, we find that the space discretization error dominates over the error on the time derivative in the tests we have performed, so that the scheme achieves high-order convergence as expected despite the use of a low-order approximation for $\partial_t s$. Second, the high-order flux f^{HO} is computed component by component. In fact since the reconstruction operators U5 and U7 ((9) and (10)) are linear, they commute with the matrices used to perform the characteristic decomposition, and there is therefore no difference in this case between component-by-component and characteristic decomposition. This contributes (along with other intrinsic differences in the formulation of the schemes) to a significant speed-up of the code (up to $\sim 50\%$, depending on the setup of the grid on the computing nodes) with respect to MP5, since there is no need to compute the system eigenvectors and apply the resulting matrix. By contrast, the MP5 reconstruction is nonlinear and does not commute with the characteristic decomposition. As a result, when using MP5 we always switch to characteristic variables, since this is known to reduce spurious numerical oscillations in high-order methods (Suresh and Huynh 1997).

Two further operations are applied on the viscosity before it is used in Eq. (15). Firstly, since the viscosity is found to be very close to zero in regions of very low rest-mass density, we improve the behaviour close to atmosphere values by simply setting the viscosity to some small and constant value ν_v at a given point x_{ijk} whenever the rest-mass density at the given point, and at all nearest neighbours, is below a certain threshold ρ_v , i.e., if

$$\rho_{i+l,j+m,k+n} < \rho_v \quad \forall l, m, n = -1, 0, 1$$

then

$$\nu_{ijk} = \nu_v.$$

In practice, therefore, (19) needs to be slightly revised and the expression for the entropy viscosity we actually implement in the numerical code is

$$\nu := \begin{cases} \nu_v & \text{if } \rho < \rho_v, \\ \min[\nu_e, \nu_{\max}] & \text{elsewhere.} \end{cases} \quad (23)$$

In all of the numerical tests presented we have used $\nu_v = 10^{-12}$ and $\rho_v = 10^{-11} M_{\odot}^{-2}$ (i.e., the threshold is 5 orders of magnitude larger than the atmosphere floor, $\rho_{\text{atmo}} = 10^{-16} M_{\odot}^{-2}$). Secondly, and following the original implementation in Guermond et al. 2011, we introduce a smoothing step which removes unwanted oscillations in the viscosity. This is accomplished by applying a five-point stencil of the form

$$\bar{\nu}_{ijk} := \sum_{l=-2}^2 \sum_{m=-2}^2 \sum_{n=-2}^2 a_l a_m a_n \nu_{i+l,j+m,k+n}, \quad (24)$$

where the coefficients a_l have values $a_0 = 0.58$, $a_{\pm 1} = 0.06$ and $a_{\pm 2} = 0.15$. This stencil in Eq. (24) is constructed so to approximate the convolution of the viscosity with a Gaussian kernel of characteristic cutoff length scale 4 times the grid spacing, in such a way that the residual of the transfer function of the target filter and of its approximation is minimised over a broad range of wavelengths (see Sagaut and Grohens 1999 for details).

In addition to dampening oscillations in the viscosity mentioned in the previous section, the necessity of the smoothing step stems from the fact that the viscosity is computed once at the beginning of every new timestep before its value is used in (15). Since the viscosity is kept constant during the series of Runge-Kutta internal steps, it ‘lags behind’ in time with respect to the solution. This issue is addressed by the smoothing procedure, but in practice we have found that this does not represent a problem in our tests. The smoothing (24) also prevents the viscosity to plunge to very small values where it should instead be

non negligible. This can happen, e.g., close to stellar surfaces as a result of oscillations in the solution. The application of the smoothing operator removes this problem by joining seamlessly the values of the viscosity in the neighbouring points.

4 Numerical tests

We report in this section the results of some of the tests obtained with the ELH method described in the previous sections. In all tests we compare the ELH results with those obtained using the monotonicity preserving, fifth-order scheme (MP5). In particular, unless otherwise stated, we couple the ELH method to the fifth-order U5 stencil (9), to make a fair and sensible comparison between methods of the same order. In few cases, however, we will also report results obtained with the seventh-order stencil U7 (10). We will refer to the corresponding schemes as to EL5 and EL7, respectively. Finally, it is useful to remark that in all of the following tests no attempt was made to tune the coefficients c_e and c_{\max} introduced in Section 3.1, and that have been set to unity here. Despite this very simple choice, the ELH method is stable and accurate in all cases considered, as the following sections make clear. At the same time, we consider it possible (if not likely) that the results could be further improved after a careful exploration of the changes in the solution upon a change of c_e and c_{\max} ; we will leave this exploration to a future work.

4.1 Special-relativistic tests

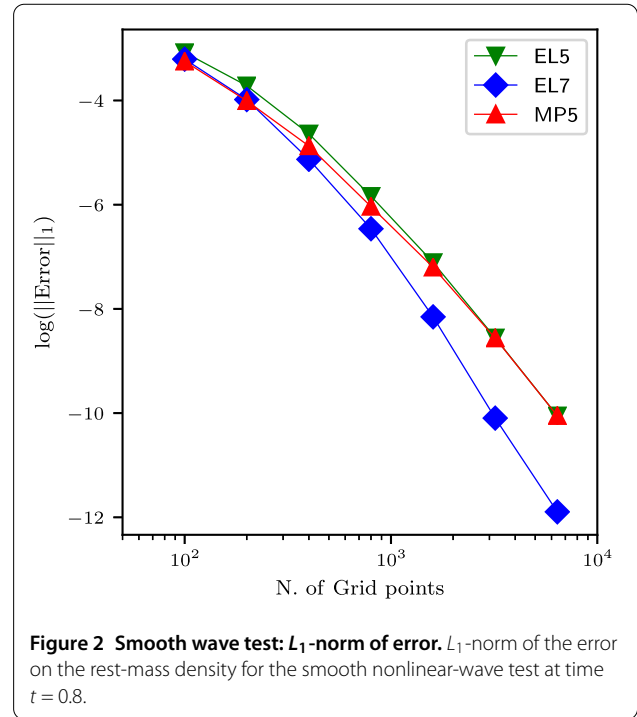
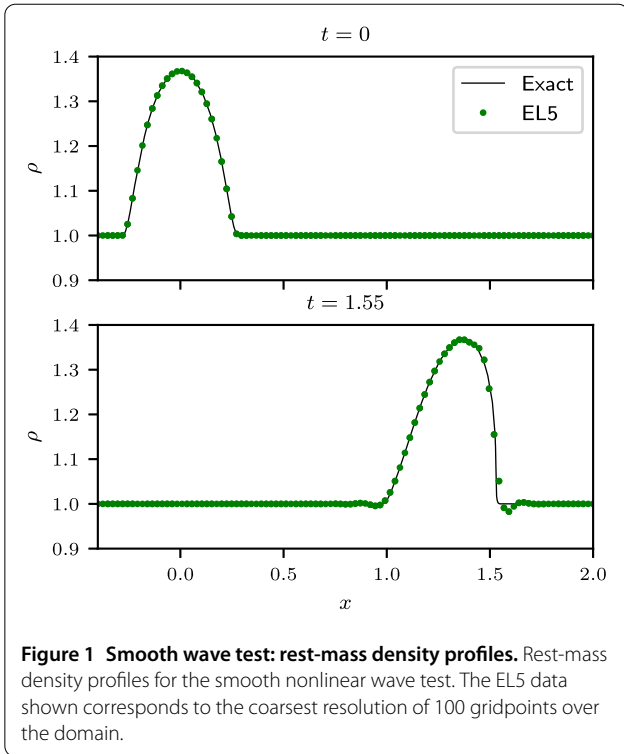
We begin with a series of mostly one-dimensional tests, performed in special-relativistic hydrodynamics, so that the metric $g_{\mu\nu}$ is fixed to the flat Minkowski metric $\eta_{\mu\nu}$ and no spacetime evolution is performed. Also, since we are mostly interested in the behaviour of the scheme in realistic astrophysical applications, we focus on just a handful of one-dimensional test cases: a smooth nonlinear wave and three shock-tube tests.

4.1.1 Smooth nonlinear wave

We first show the accuracy of the scheme in the case of a smooth solution and measure rigorously its convergence order so as to show that the entropy-driven limiter does not spoil the convergence properties of the high-order method upon which it is built. This test has been discussed in Radice and Rezzolla (2012) (which adapted it from Zhang and MacFadyen 2006). In short, we consider a one-dimensional, large-amplitude, smooth, nonlinear wave with initial rest-mass density profile given by

$$\rho_0(x) = \begin{cases} 1 + \exp[-1/(1 - x^2/L^2)] & \text{if } |x| < 1, \\ 1 & \text{elsewhere,} \end{cases} \quad (25)$$

where $L = 0.3$. The initial data employs a polytropic EOS, $p = K\rho^{\tilde{\gamma}}$, with $K = 100$ and $\tilde{\gamma} = 5/3$, and we then evolve



it with the ideal-fluid EOS (4) with $\Gamma = 5/3$. Since in this test discontinuities are absent (so that $\tilde{\gamma} = \Gamma$) and there are no stability issues, we use as time integrator the standard fourth-order RK4 method with a timestep of ~ 0.13 times the grid spacing.

The analytic solution of the test (shown in Figure 1 with a black solid line) is represented by a wave profile propagating towards the right and ‘steepening’ in the direction of its motion. At time $t_c \simeq 1.6$, the wave develops a shock, leading to a sharp discontinuity. Up to the formation of the caustic, the analytic solution can be computed using the method of characteristics (Anile 1990) on a Lagrangian grid. We obtain an accurate enough approximation by computing it on a very fine grid of 10^5 gridpoints and interpolating the solution using cubic splines on the Eulerian grid. This solution, which we refer to as the ‘exact’ solution, is then used as the reference against which the numerical solutions are compared.

We perform this test with both the EL5 and EL7 schemes of the ELH method to validate that high-order schemes can be employed with great ease in our approach by simply swapping a lower-order stencil for a higher-order one; this operation is far more demanding in standard finite-volume HRSC schemes.

Figure 2 shows the L_1 -norm of the error with respect to the analytic solution at time $t = 0.8$ for the various schemes and at various resolutions. The latter are parametrized by the number of gridpoints used on the x axis, and we

have considered seven different resolutions, each twice as fine as the preceding one, going from 100 gridpoints up to 6,400. The different lines in Figure 2 show that at the lowest resolutions all schemes show very similar errors, MP5 being the most accurate by a small margin. As the resolution is increased, however, the gap in accuracy between EL5 and MP5 decreases and disappears at very high resolutions. The error curve of EL7, being a higher-order scheme, decreases much more rapidly with resolution, so that its error at the highest resolution of 6,400 gridpoints is two orders of magnitude lower than for the fifth-order schemes.

We also compute the convergence order of the various schemes using the data at resolutions of 1,600 and 3,200 gridpoints and after comparing it with the ‘exact’ solution. The result is shown in Figure 3 as a function of time to the development of the shock. The computed order should be equal to the nominal order of each scheme as long as the solution is smooth, gradually degrading to first order as the caustic is approached. Every scheme matches this description, in particular EL5, whose convergence order is almost exactly five. EL7 similarly appears to saturate just below its nominal convergence order of seven. Deviations from the nominal convergence order of each scheme are due to contaminations from other error sources, which become increasingly significant at high resolution; these are: the truncation error due to the time-integrator, the accuracy of the inversion from conservative to primitive variables,

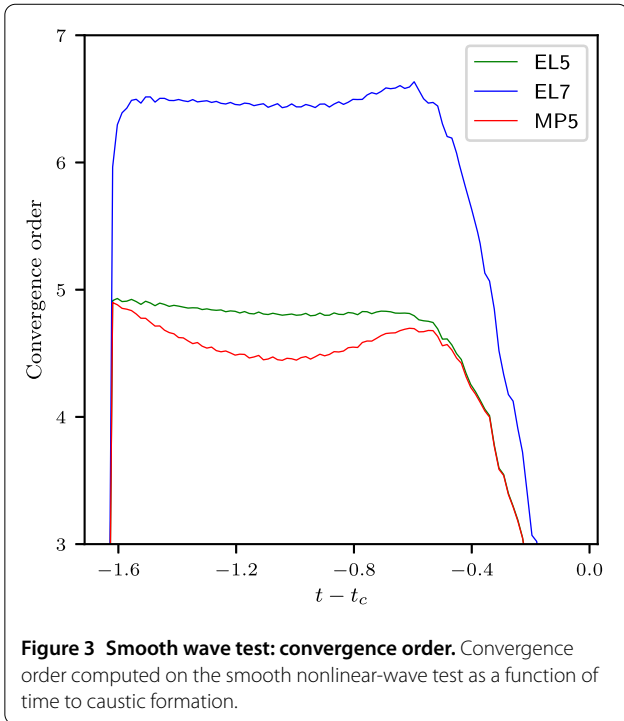


Figure 3 Smooth wave test: convergence order. Convergence order computed on the smooth nonlinear-wave test as a function of time to caustic formation.

or the low-order approximation for the evolution of the entropy (cf., Eq. (22)). What is relevant here is that the EL method does not interfere with the convergence properties of the underlying stencil and can exploit their accuracy.

4.1.2 Shock-tube tests

We choose as a first shock-tube test the special-relativistic version of the classical Sod test (Sod 1978). In this case, the adiabatic index for both the polytropic initial data EOS and the ideal-fluid evolution EOS is $\Gamma = 1.4$ and the right (R) and left (L) initial states are

$$\begin{aligned}
 (\rho_R, v_R, p_R) &= (0.125, 0, 0.1), \\
 (\rho_L, v_L, p_L) &= (1, 0, 1).
 \end{aligned}
 \tag{26}$$

The analytic solution consists in a left-going rarefaction wave and a right-going shock wave separated by a right-going contact discontinuity. We perform the test with a variety of spatial resolutions ranging from $\Delta x = 0.01$ to $\Delta x = 3.125 \times 10^{-4}$, and a timestep $\Delta t = 0.1 \Delta x$.

Figure 4 shows the test results at time $t = 0.6$ for both the EL5 and MP5 schemes at resolution $\Delta x = 1.25 \times 10^{-3}$. Both schemes capture the main features of the solution, with the shocks being captured within ~ 3 gridpoints, as are the constant states in the pressure and velocity. However, the EL5 scheme displays some oscillations downstream of the shock as well as over- and undershoots around the location of the discontinuities and in the transition between the rarefaction wave and the surrounding flat regions, while the MP5 scheme is able to resolve the solution avoiding such artefacts. This is not surprising since MP5 is a monotonicity preserving scheme (the number of local maxima and minima cannot increase by effect of this method, therefore over- and undershoots cannot occur by construction)




Figure 4 Sod test: solution profiles. Profiles of the rest-mass density (left), velocity (center) and pressure (right) for the special-relativistic Sod test at $t = 0.6$. The solution is computed on a grid of 800 points. The EL5 scheme correctly captures the features of the solution despite oscillations at the discontinuities.




Figure 5 Sod test: solution at different resolutions. Rest-mass density profiles, zoomed on the right edge of the rarefaction wave (left) and on the shock (right) for the special-relativistic Sod test at $t = 0.6$, computed with EL5 at different resolutions (parametrized by the number of points on the x axis). The oscillations in the solution can be seen converging away with resolution.




Figure 6 Sod test: viscosity at different resolutions. Profiles of the viscosity in logarithmic scale for the special-relativistic Sod test at $t = 0.6$, computed with EL5 at different resolutions (parametrized by the number of points on the x axis). The four peaks correspond to the four different features of the solution, i.e., from left to right, the edges of the rarefaction wave, the contact discontinuity and the shock (vertical dashed lines highlight their location). As the resolution increases, they tend to delta functions.

while EL5 is not. We should remark, however, that this property of MP5 is valid only for scalar equations in one spatial dimension, and it basically accounts for the differences in the behaviour of the two schemes. It should also be stressed that the EL5 scheme is indeed stable and that the oscillations that are present in the solution converge away with resolution (see Figure 5).

The behaviour of the viscosity is displayed in Figure 6. It presents four well distinct peaks, each corresponding to the four nonlinear waves generated by the Riemann problem and corresponding to the edges of the rarefaction wave, where the solution is continuous but non-smooth, of the contact discontinuity and of the shock. The viscosity is higher in correspondence with the latter, and decreases by several orders of magnitude for the other three. It can also be seen clearly how the peaks in the viscosity sharpen as the resolution is increased, mirroring the decreasing size of the aforementioned features (see Figure 5), and seemingly tending towards a delta function at infinite resolution, as expected.

The second shock-tube test we select is a more extreme ‘blast-wave’ test (Martí and Müller 2003). In this case, the adiabatic index used is $\Gamma = 5/3$ and the right and left initial states are

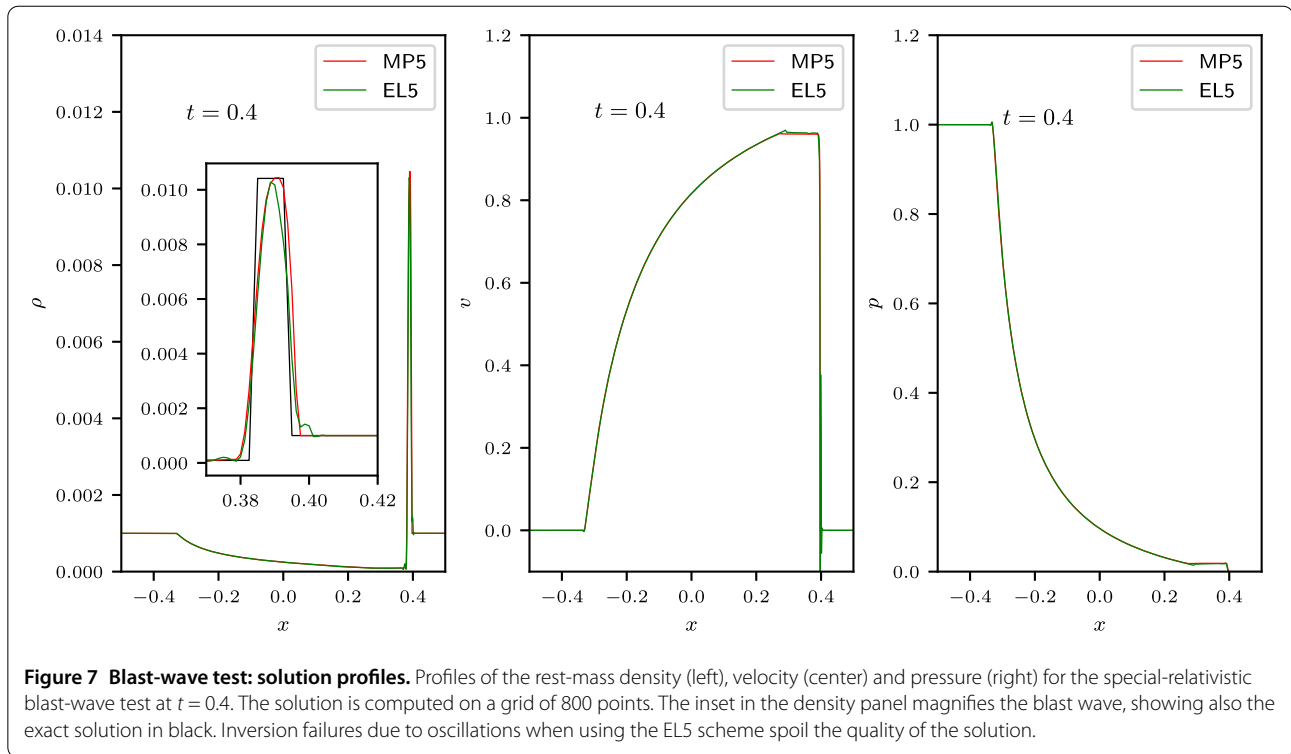
$$\begin{aligned} (\rho_R, v_R, p_R) &= (10^{-3}, 0, 1), \\ (\rho_L, v_L, p_L) &= (10^{-3}, 0, 10^{-5}). \end{aligned} \tag{27}$$

The exact solution consists in a right-going shock wave, followed by a contact discontinuity and a left-going rarefaction wave. We employ the same resolutions and time-step choices as for the Sod test.

Figure 7 is similar to Figure 4, but relative to the blast-wave test at time $t = 0.4$. We should remark that this is a very extreme test (the pressure has an initial jump of five orders of magnitude) in which the contact discontinuity and shock wave move at essentially the same speed, yielding a very narrow constant rest-mass density state between the two. The oscillations in the EL5 scheme data are in this case more severe than in the Sod shock-tube test, especially around the shock location. As a result, the solution with the EL5 scheme tends to a general decrease of the pressure between the rarefaction wave and the shock wave, whose relative value is however of $\lesssim 7\%$ at most; the MP5 scheme performs better and has a relative error in pressure that is $\sim 1\%$. In both cases, the agreement with the exact solution improves with resolution.

Finally, we perform a three-dimensional shock-tube problem, involving non-grid-aligned shocks i.e., the relativistic-explosion test. The initial data in this case is given by

$$\begin{cases} (\rho, v_i, p) = (1, 0, 1) & \text{if } r \leq 0.4, \\ (\rho, v_i, p) = (0.125, 0, 0.1) & \text{otherwise,} \end{cases}$$



where r is the distance from the origin. The computational domain is a cube of side 1 centered on the origin, and we use a grid spacing of $\Delta x = 0.01$ and a timestep $\Delta t = 0.1 \Delta x$. The adiabatic index for this test is again $\Gamma = 1.4$. The features of the solution are similar to those of the Sod test i.e., an ingoing rarefaction wave and an outgoing shock, separated by an outgoing contact discontinuity. Note however that because of the spherical symmetry of the test (compared to the planar symmetry in the Sod case), the regions at the two sides of the contact discontinuity are no longer constant states in rest-mass density, velocity and pressure, but display a smooth radial dependence.

Figure 8 shows the rest-mass density for this test at time $t = 0.25$, on the (x, y) plane as well as on the x axis. Both EL5 and MP5 perform very similarly, with differences being barely noticeable in the two-dimensional plot. The curves on the x axis reveal that while both schemes capture the features of the solution, as in the Sod test, the EL5 scheme is slightly more oscillatory.

Overall, these shock-tube tests demonstrate how the entropy-driven hybridisation of the high-order stencil is sufficient to stabilise the scheme even for discontinuous initial data and it is remarkable that such a simple scheme can achieve good accuracy.

4.2 Three-dimensional general-relativistic tests: neutron stars

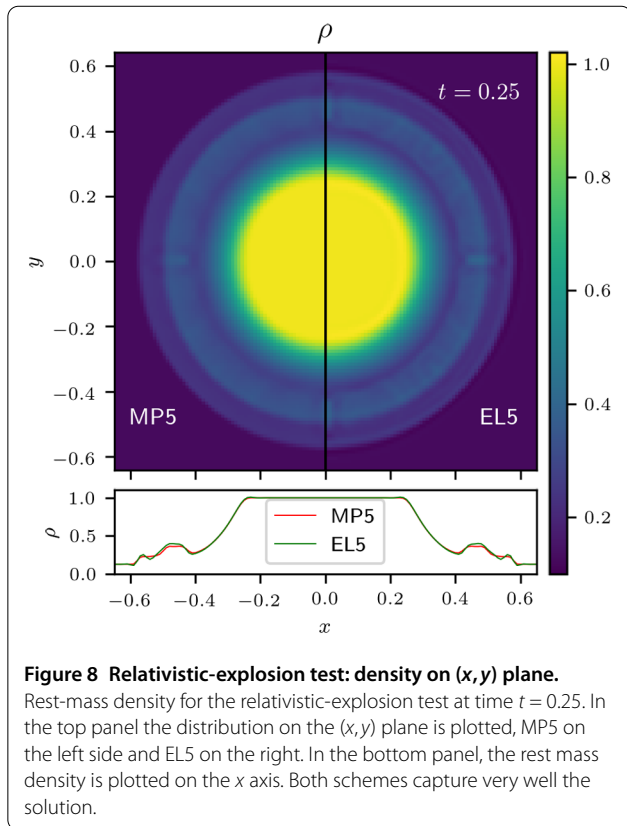
We next test the EL5 scheme against a series of three-dimensional tests mostly based on the evolution of single,

isolated neutron stars in general relativity (with the exception of grazing-collision test of Section 4.2.6). In each test we employ for the evolution the ideal-fluid EOS (4) with $\Gamma = 2$. The neutron star initial data is constructed using a polytropic EOS $p = K\rho^{\tilde{\gamma}}$ also with $\tilde{\gamma} = 2$ and $K = 100 M_{\odot}^{-2}$.

4.2.1 Isolated star in the Cowling approximation

The first test we perform is the evolution of a stable non-rotating (or TOV, from Tolmann-Oppenheimer-Volkoff) neutron star in a fixed spacetime (i.e., adopting the Cowling approximation) with the goal of assessing the properties of the EL5 scheme over long timescales. Despite its conceptual simplicity (a TOV is just a static solution of the Einstein-Euler equations) the test can be rather challenging. This is because in this test the location of the stellar surface, which is the hardest feature to simulate due to the steep gradient in the hydrodynamics variables, is essentially stationary; as a result, errors can accumulate and grow, affecting the accuracy of the simulation. This behaviour is to be contrasted with the typical situation encountered when evolving inspiralling binary neutron stars, where the stellar surfaces move very supersonically with respect to the floor and most of the errors at the surface are absorbed into the shocks.

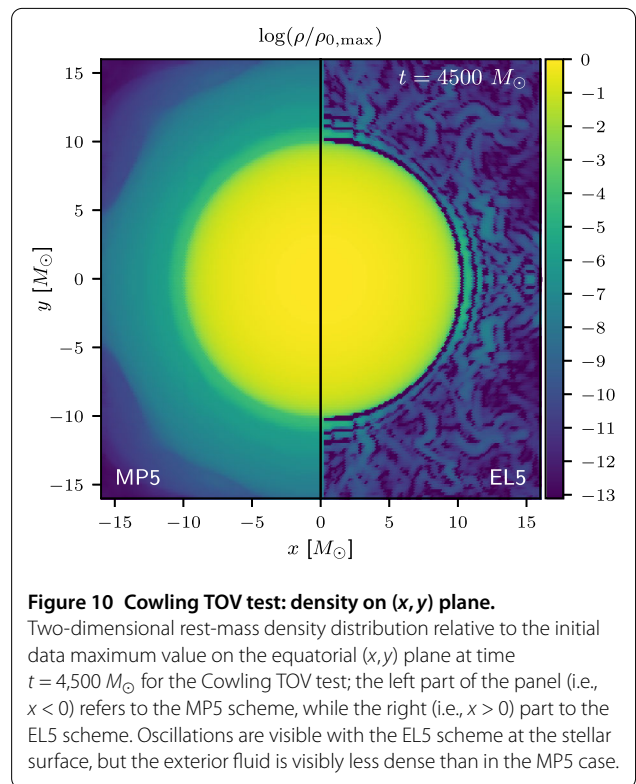
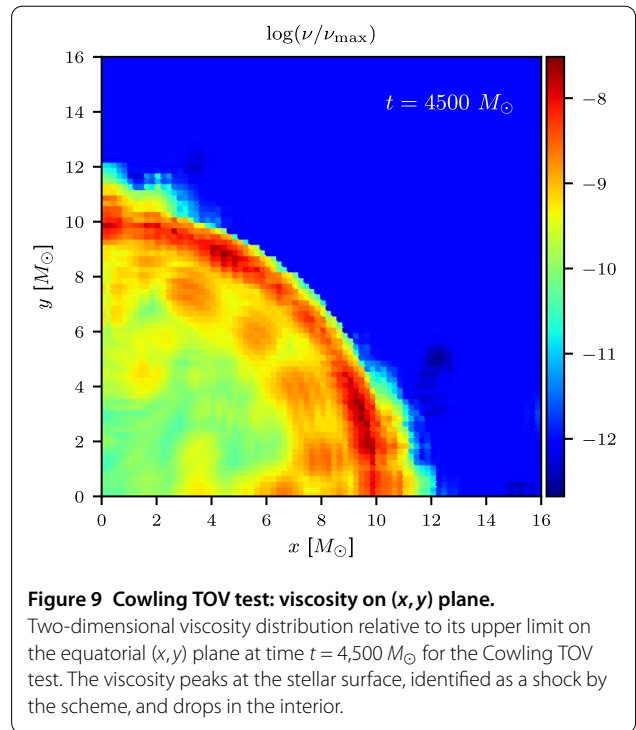
For this test we build and evolve a TOV model with central rest-mass density $1.28 \times 10^{-3} M_{\odot}^{-2}$, yielding a (baryon) rest mass of $1.5 M_{\odot}$ and a radius of $\sim 10 M_{\odot}$. We perform the test on a single refinement level with outer boundaries



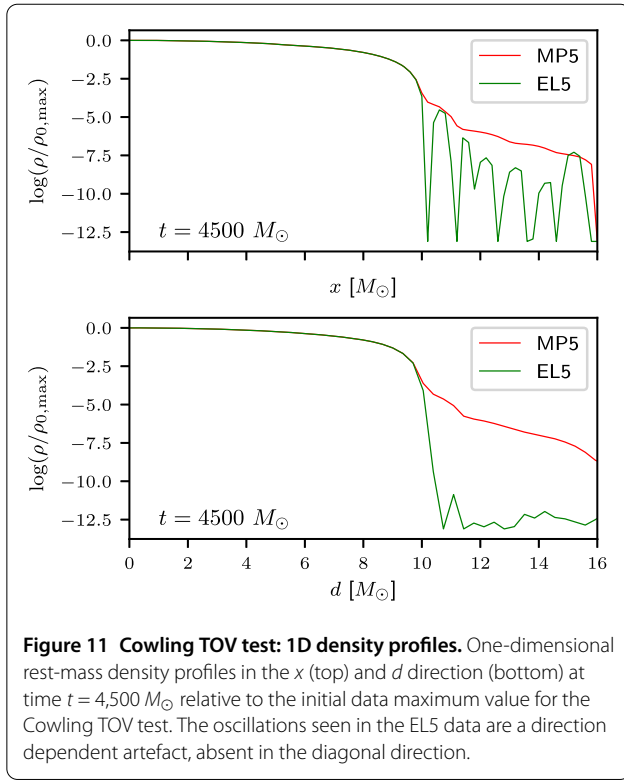
placed at $16 M_{\odot}$ and a resolution of $\Delta^i = 0.2 M_{\odot} \simeq 0.3$ km. The timestep is set to 0.15 times the grid spacing, and the time integrator is RK3.

Figure 9 reports the distribution of the viscosity on the equatorial plane, which clearly shows a local annular peak around the location of the stellar surface, where the hydrodynamical variables experience the most violent variations, leading to large values of the viscosity. In the external low-density fluid, the viscosity is set to a small constant value almost everywhere as detailed in Section 3.2. The inner part of the neutron star is expected to be isentropic, as it consists of a shock-free perfect fluid. Indeed, in the stellar interior the viscosity is nonzero but also 10^2 to 10^3 times smaller than at the surface and does not significantly affect the evolution. Quite generally, these features of the viscosity profile are typical in all the tests we considered, whenever a sharp matter/vacuum interface is present.

The general behaviour of the EL5 scheme, when compared to the MP5 scheme, is well illustrated by Figure 10, showing the rest-mass density distribution on the equatorial plane for the two schemes (the left part of the panel, i.e., for $x < 0$, refers to the MP5 scheme, while the right part, i.e., for $x > 0$, to the EL5 scheme^c). As it can be seen from the figure, both schemes accurately capture the solution in the stellar interior, but significant differences arise at the surface and in the exterior. The MP5 scheme shows a rather diffusive behaviour, with a smooth transition to



the external ‘vacuum’ (i.e., to a region close to the rest-mass density floor) and extended low-density tails. The EL5 scheme, on the other hand, produces a sharper edge.



Oscillations in the solution can be seen just outside of the star, resulting in shell-like structures around the surface, which are particularly noticeable in the coordinate axes directions. The stellar exterior is much closer to the vacuum with the EL5 scheme and, in contrast to MP5, it also displays small-scale dynamics at very low densities.

The properties of the oscillations present in the solution computed with the EL5 scheme are made clearer in Figure 11, which shows the rest-mass density profiles along different radial cuts. Along the x direction, the oscillations in the EL5 data have large amplitude and a similar behaviour is observed along the y and z axes. On the other hand, on the three-dimensional diagonal (i.e., along the $x = y = z$ line), the EL5 scheme manages to capture the sharp transition between the stellar interior and the outside vacuum almost perfectly, without significant oscillations or other artefacts. By contrast, the use of the MP5 scheme leads to smooth, rest-density profiles that are only slowly decaying in all directions.^d

The direction-dependent behaviour shown in Figure 11 for the EL5 scheme is due to the well-known anisotropy of the phase error common to finite-differencing schemes (Vichnevetsky and Bowles 1982; Lele 1992). The MP5 scheme is able to mask this behaviour, but at the price of sacrificing the ability to sharply define stellar surface. We expect that the performance of the EL5 scheme could be improved through the use of multidimensional stencils (i.e., employing a multidimensional interpolation in the re-

construction step), as opposed to the current approach in which the stencil is simply oriented in the direction of the flux to be reconstructed.

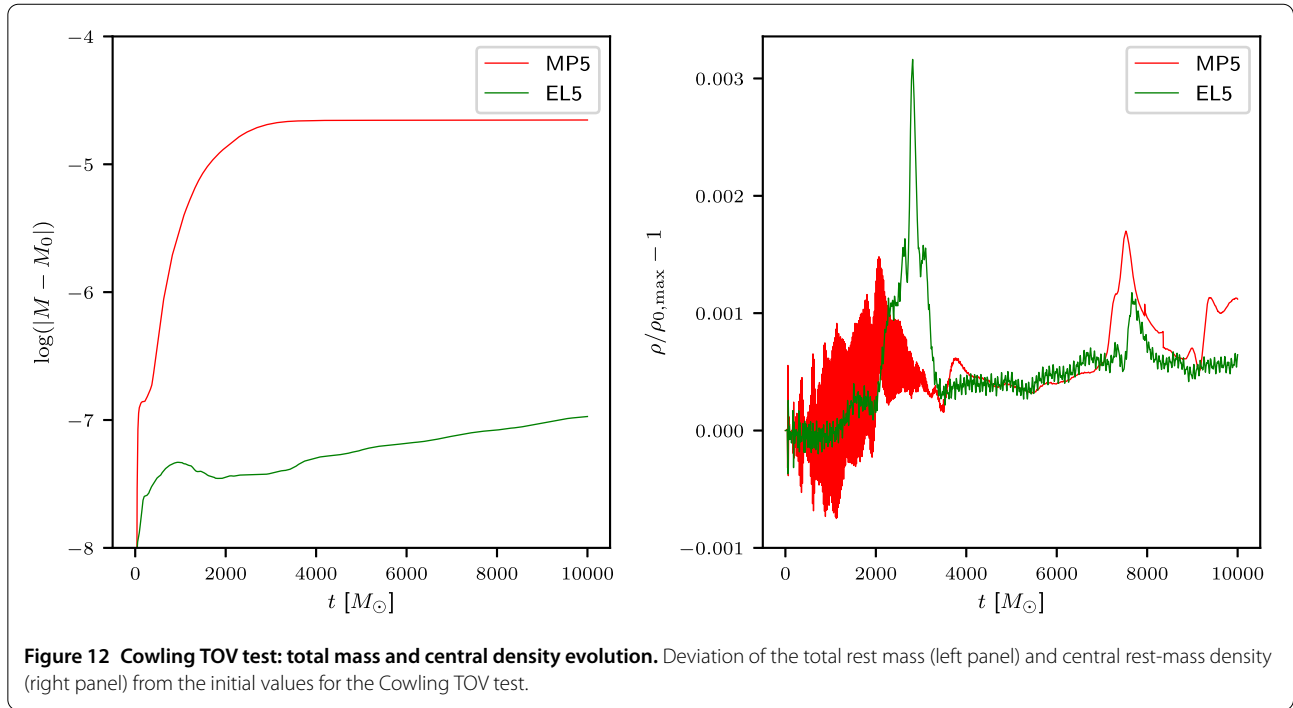
The quantitative differences between the two schemes are better captured in Figure 12, where the evolution of the total rest mass and of the central rest-mass density are shown. We recall that the total rest mass (or baryon mass), is defined as

$$M := \int \rho W \sqrt{\gamma} d^3x, \quad (28)$$

where the integral is performed over the whole computational domain. From the continuity equation (35) follows that it should be conserved in absence of a net total flow of matter in or out of the domain. The numerical schemes we employ are conservative (see e.g., Leveque 1992), and therefore preserve the value of the rest mass to the one determined by the initial data. Nonetheless, violations of this conservation can take place in at least three different ways. First, winds originating at the stellar surface (physically, as e.g., in binary neutron star merger, or spuriously as in a stationary case such as the present one) can yield a net loss of mass when they reach the outer boundary and leave the computational domain. Second, matter can be spuriously created or destroyed, in a way that is hard to control, because of floating-point or interpolation errors at the boundaries of refinement levels (this is not the case for this particular test clearly, since we employ a single grid, but it is relevant for the following ones). Finally, when a value of the density is floored mass is spuriously created or destroyed. It is therefore important to characterize the interplay between the numerical scheme and these grid related effects.

The left panel of Figure 12 shows deviations, in absolute value, of the rest mass from the initial value for the two schemes. The EL5 scheme is evidently much better at conserving mass in this test than MP5, leading to a cumulative deviation of $\sim 10^{-7} M_{\odot}$ which is almost three orders of magnitude smaller than the MP5 value.

The central rest-mass density also undergoes an evolution (right panel of Figure 12), with oscillations triggered by the treatment of the stellar surface. Both schemes perform at a similar level of accuracy, with relative variations from the initial value no greater than about 0.3% (even though spurious peaks are present in both data series at various times). The short term behaviour of the two schemes is noticeably different, and the frequency content in the two data series appears different, with the MP5 scheme seeming to show more pronounced high-frequency modes. However, at later times both schemes appear to relax and oscillations decrease significantly in amplitude.



4.2.2 Isolated star in a dynamical spacetime

We then proceed to relax the Cowling approximation and test the entropy-limited method coupled with a dynamically evolved spacetime. As first step, we evolve the same initial data for a isolated stable star as in the previous section (i.e., with central density $1.28 \times 10^{-3} M_{\odot}^{-2}$, baryon mass of $1.5 M_{\odot}$ and radius $\sim 10 M_{\odot}$). We perform the test on a grid consisting of three refinement levels centered on the star with sides lengths 16, 32 and $60 M_{\odot}$ from finest to coarsest, and with a constant refinement factor of 2. The spatial resolution of the innermost and finest level is set to $\Delta^i = 0.2 M_{\odot} \simeq 0.3 \text{ km}$, and the timestep to 0.15 times the grid spacing. This factor is largest possible to guarantee the positivity of the rest-mass density (see discussion in Section 3.1 and Radice et al. 2014a for details). The atmosphere value of the density is set to $10^{-16} M_{\odot}^{-2}$, that is, almost 13 orders of magnitude smaller than the maximum value. As a time integrator we select the third-order SSP Runge-Kutta RK3. Unless stated differently, we employ the same grid setup for each one of the following single star tests.

Figure 13 shows the distribution of rest-mass density on the equatorial plane for this test, again with the MP5 and EL5 schemes shown on the left and right parts of the panel, respectively. It can be appreciated how the MP5 scheme produces rest-mass tails which are even more dense and extended than in the Cowling case, making the near vacuum solution obtained by the EL5 scheme all the more striking.


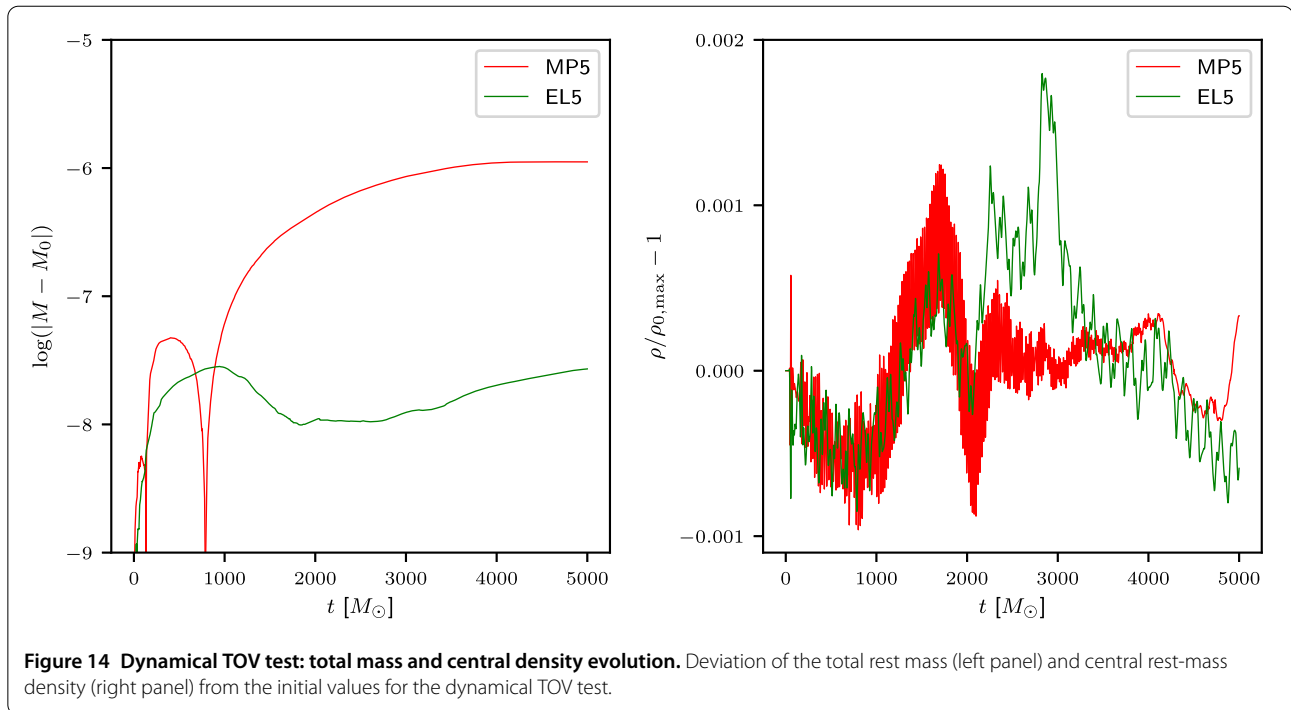


Figure 13 Dynamical TOV test: density on (x, y) plane. Two-dimensional rest-mass density distribution relative to the initial data maximum value on the equatorial (x, y) plane at time $t = 4,485 M_{\odot}$ for the dynamical TOV test. The matter tails are even more extended in MP5 case compared with the Cowling test, EL5 instead preserves its behaviour at the stellar surface and exterior.

Another difference from the Cowling test can be seen in the conservation of the rest mass (left panel of Figure 14). In this case too, the EL5 scheme is able to conserve the initial value to an accuracy roughly two orders of magnitude

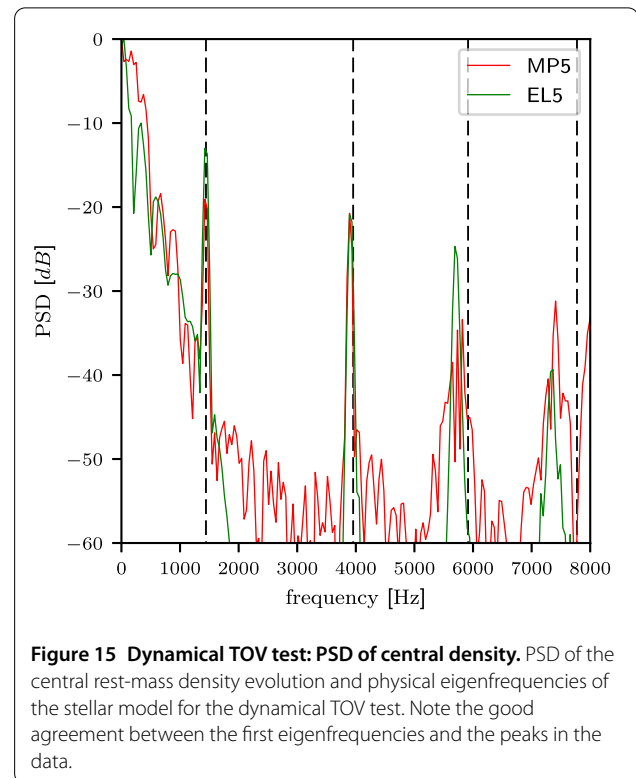


better than the MP5 scheme. Furthermore, it is interesting to notice how the behaviour of the EL5 scheme is much more smooth and predictable; MP5 by contrast displays both spurious losses and gains of mass, which lead to the zero crossings clearly visible in the figure. This is due to interpolation errors arising during the restriction and prolongation operations between different refinement levels. These errors are more severe with MP5 due to the presence of long tails of low density matter in the stellar exterior, as we checked by varying the extent of the refinement levels. In contrast, the EL5 scheme is less affected since the exterior of the star (especially away from the coordinate axes) is nearly vacuum.

The evolution of the central rest-mass density, as shown in the right panel of Figure 14, is similar to the one shown in the previous section for the Cowling approximation, with both schemes varying no more than 0.2% from the initial value, but with MP5 displaying oscillations at much higher frequency.

To further investigate this point, we compute the power spectral density (PSD) of the density evolution, in order to quantitatively gauge the differences between the two schemes. The PSD is computed over the first 5,000 M_{\odot} of data and with the use of a Hann window function. Before computing the PSD, any linearly growing component of the signal is removed via a least-squares fit.

The PSDs for both schemes are shown in Figure 15 along with the oscillation frequencies of this stellar model as computed perturbatively following the methods discussed in Yoshida and Eriguchi (2001), Takami et al. (2011), and



shown as vertical dashed lines. For both schemes, the PSD is dominated by a low-frequency component due the well-known secular changes in the central rest-mass density (Font et al. 2002) and that disappear with resolution.

However, peaks are clearly visible above the noise. The lowest-frequency peaks correspond to the fundamental oscillation mode of the star and its first overtone, while the following ones are higher overtones and are progressively more offset from the corresponding perturbative eigenfrequencies. The peaks in the EL5 data appear to be more clearly identifiable and less broad than in the MP5 case. Above $\sim 8,000$ Hz, and not shown in Figure 15, the MP5 scheme shows significant high-frequency components, clearly visible in the first part of the corresponding curve in Figure 14, but with a rather disordered spectrum. These same frequencies are instead greatly suppressed in the EL5 scheme. Overall, also this test highlights that the EL5 scheme captures quite well the physical behaviour of the system as expected from perturbative methods and is free from some of the artefacts which appear instead in the evolution with the MP5 scheme.

4.2.3 Perturbed isolated star

The next test we perform is a slight modification of previous setup, i.e., we evolve the same isolated neutron star model, but applying a small velocity perturbation to the initial solution. The perturbation consists of a radially outgoing velocity growing linearly in radius to a maximum value of 0.005.

We employ this scenario, more realistic than the simple smooth-wave test of Section 4.1.1, to measure the convergence order of the EL5 and MP5 methods. We performed three sets of simulations at resolutions 0.24, 0.12 and $0.06 M_{\odot}$ on the finest level, extracting the evolution of the rest-mass density over time from each one. The initial perturbation is added so that the density evolution is not dominated by the truncation error, but displays a cleaner behaviour. Otherwise, as the resolution is increased, the density evolution would show additional high-frequency modes, which would make the dependence on resolution discontinuous, making it difficult to compute the instantaneous convergence order.

Using the values of the L_1 -norm of the rest-mass density over the domain at the three resolutions we also computed the instantaneous convergence order M_{\odot} as shown in Figure 16. Because this is the instantaneous convergence order and because the underlying system is oscillating, the curves are somewhat noisy (especially for MP5); however, when taking the running average, both schemes generally show a convergence order just below three, consistent with the results in Radice et al. (2014a, 2014b). It is also however apparent how EL5 maintains a fairly constant order of convergence through time, while the behaviour of MP5 is more irregular, especially at later times.

While the hydrodynamics schemes are both formally fifth-order accurate, other parts of the algorithm operate at different degrees of accuracy. In particular, the time integrator is third-order accurate, which most likely accounts

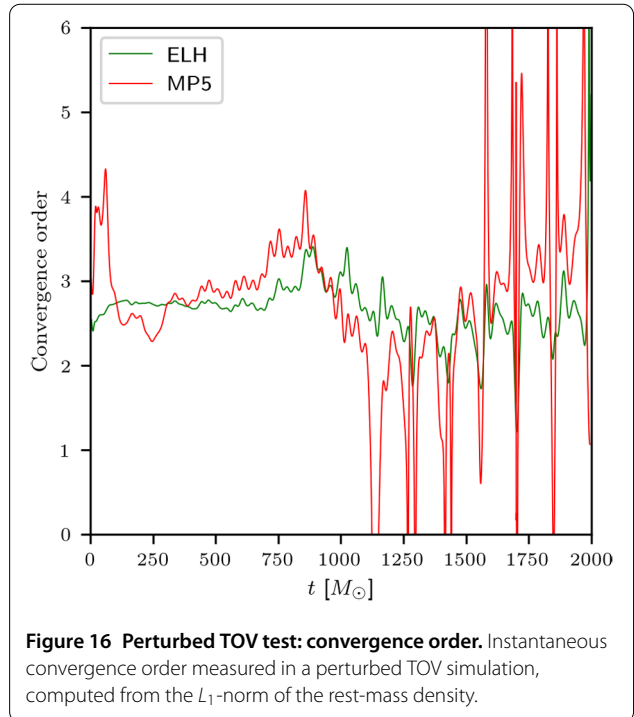


Figure 16 Perturbed TOV test: convergence order. Instantaneous convergence order measured in a perturbed TOV simulation, computed from the L_1 -norm of the rest-mass density.

for the convergence order being closer to three than to five. The result is also consistent with the ones found for the MP5 scheme in Radice et al. (2014a, 2014b), Radice et al. (2015). Overall, this test highlights how both the ELH and MP5 schemes perform fairly consistently over time, with no major loss of accuracy.

4.2.4 Migration test

Another important test in our series is the migration of a TOV star moving from a solution on the unstable branch of equilibrium solutions to a stable one. We recall that for any given EOS, increasingly massive but stable TOV models can be constructed by considering increasingly large values of the central rest-mass density. This can continue until a maximum mass is reached, at which point, an increase of the central rest-mass density corresponds to a decrease of the mass of the star. Models on this second branch of the mass/central-rest-mass density curve are unstable, and if a perturbation is present will evolve to either a stable configuration or collapse to a black hole. This is precisely the physical scenario that the migration test simulates: we construct a model on the unstable branch of the mass/density curve and force its ‘migration’ to a stable configuration by applying a suitable velocity perturbation.

This is a common test for numerical relativity codes (see, e.g., Font et al. 2002; Baiotti et al. 2003, 2005; Cordero-Carrión et al. 2009; Thierfelder et al. 2011), and has been studied in detail in Radice et al. (2010). In particular, we build a nonrotating stellar model on the unstable branch of the equilibrium solutions and with central rest-mass density of $7 \times 10^{-3} M_{\odot}^{-2}$ (yielding a total rest mass of $1.6 M_{\odot}$




Figure 17 Migration test: density on (x,y) plane. Two-dimensional rest-mass density distribution relative to the initial data maximum value on the equatorial (x,y) plane at time $t = 1,167 M_{\odot}$ for the migration test. Virtually no difference can be detected in the two schemes behaviour.

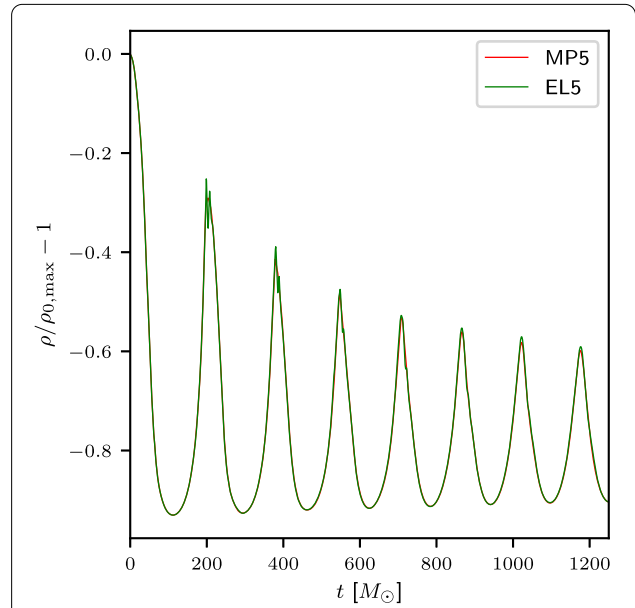


Figure 18 Migration test: central density evolution. Central rest-mass density in the migration test. The agreement between the two schemes is apparent over the whole evolution, apart from high-frequency modes at the maxima in EL5 data.

and a radius of $6 M_{\odot}$). The migration is then triggered by injecting a radially outgoing velocity perturbation where the velocity grows linearly in radius, reaching a maximum value of 0.01. The star then undergoes a series of violent expansions and contractions as it migrates to the stable branch and then settles on the new equilibrium. During each contraction and expansion strong shocks are formed, and the shocked matter is ejected at large velocities.

In Figure 17 we show the rest-mass density distribution on the equatorial plane for both schemes and during one of the contractions of the star, just before the central rest-mass density reaches a maximum (cf., Figure 18). The snapshot clearly shows that both the EL5 and MP5 schemes produce almost identical results for this test. This is not surprising and mainly due to the matter outflow driven by the stellar oscillations, which rapidly fills the domain and removes the sharp feature of the stellar surface, which is the most problematic structure to resolve and the main difference in the two schemes.

The evolutions of the maximum rest-mass density are shown in Figure 18, where they are reported as normalized to the initial value. The agreement between the two schemes is extremely good during the entire evolution and the main difference between the two solutions is the presence of some high-frequency modes near the maxima of the density in the EL5 data. Such oscillations are the result of inward-propagating shock waves generated in the outer layers of the star during the contraction phase. Figure 19 shows a magnification of the behaviour of the maximum




Figure 19 Migration test: central density evolution at different resolutions. Magnification of the central rest-mass density evolution around the first contraction of the migrating star. The low-resolution data (thin curves) corresponds to Figure 18 and to a grid spacing of $0.2 M_{\odot}$. The high-resolution data (thick curves) corresponds to a grid spacing of $0.086 M_{\odot}$. The high-frequency modes present in the EL5 data at low resolution persist at high resolution in both schemes.

rest-mass density at the peak of the first contraction, comparing not only the two schemes but also the evolutions with two different resolutions. At high resolution, both the MP5 and the EL5 scheme show small-scale and high-frequency oscillations that are less pronounced in the low-resolution data. Interestingly, these oscillations are essentially smoothed out in the low-resolution run of the MP5 scheme, while they are very visible in the low-resolution EL5 run. This seems to indicate that the two schemes tend, with increasing resolution, towards a solution where the small-scale oscillations are present and therefore physically correct and not a numerical artefact. Finally, as the evolution progresses, the contraction/expansion phases become less and less violent as part of the kinetic energy is converted into internal energy, thereby leading to milder and milder shocks, and the high-frequency oscillations in the central rest-mass density all but disappear.

4.2.5 Isolated rotating neutron star

As the last test case for a stable (or metastable) isolated relativistic stars we consider the evolution of a rapidly and uniformly rotating star. More precisely, we set up axisymmetric initial data relative to a uniformly rotating neutron star governed by a polytropic EOS with $K = 100 M_{\odot}^{-2}$ and $\Gamma = 2$, having a central rest-mass density of $1.28 \times 10^{-3} M_{\odot}^{-2}$ and a polar to equatorial axis ratio of 0.8 using the RNS code (Stergioulas and Friedman 1995). This results in a star with total rest mass $1.6 M_{\odot}$, radius $10 M_{\odot}$, rotation frequency $f = 673.2$ Hz (about 60% of the mass shedding frequency) and dimensionless angular momentum $J/M^2 = 0.46$. Also in this case the spacetime is evolved in time despite the solution being stationary.

In Figure 20 we report again the rest-mass density distribution on the equatorial plane for both schemes and at time $t = 4,300 M_{\odot}$, that is, after about 14 rotation periods. The figure clearly illustrates that that both schemes evolve the rotating star with no noticeable problems and that, as already seen in the case of nonrotating stars, the part of the domain exterior to the stellar surface rapidly fills with matter. Also in this case, the behaviour of the two methods in the low-density regions is rather different, with the MP5 scheme yielding to a volume which is filled by uniform but comparatively higher-density material, while the EL5 scheme produces a stellar exterior which has lower-density matter but with small-scale condensations (cf., Figures 10 and 13 for the equivalent behaviour in the absence of rotation).

However, in contrast with the behaviour seen in the case of nonrotating stars, the dynamics of the low-density material in the stellar exterior results in a degradation of the conservation of mass for the EL5 scheme, as shown in the left panel of Figure 21 (cf., the left panels of Figures 12 and 14 for the equivalent behaviour in the absence of rotation). The deviation of the total rest-mass density from its initial

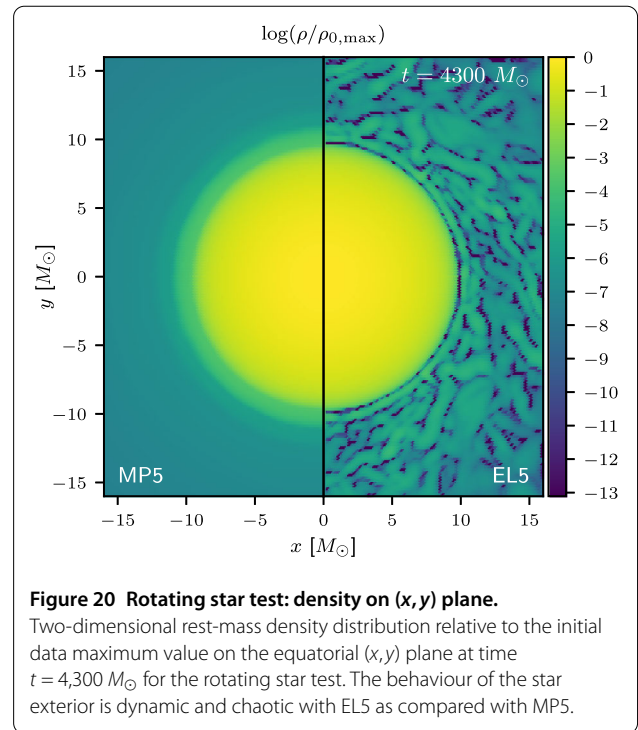



Figure 20 Rotating star test: density on (x, y) plane.

Two-dimensional rest-mass density distribution relative to the initial data maximum value on the equatorial (x, y) plane at time $t = 4,300 M_{\odot}$ for the rotating star test. The behaviour of the star exterior is dynamic and chaotic with EL5 as compared with MP5.

value is more than one order of magnitude larger for the EL5 scheme than for MP5 and reaches values of $\sim 10^{-5} M_{\odot}$. This is the result of failures in the conversion from the conserved variables to the primitive ones, triggered by oscillations in the solution: the solution can be evolved to an unphysical state, and in this case the rest-mass density could reach values below the atmosphere floor value; if so, the conversion routine resets the affected cells to the atmosphere value, thereby spuriously creating mass. We speculate that most of these failures result from the large tangential velocity that is acquired by the shell-like distribution of matter that builds up in the case of the EL5 scheme and that is present already in the nonrotating case. While rather innocuous in the absence of rotation, this shell of matter can fling material to large distances (but within the computational domain) and lead to a much more chaotic dynamics of the fluid in the low-density regions (see the discussion in Section 3.2.3 of Radice et al. 2014a).

To assess the impact of the fluid dynamics in the stellar exterior we report the evolution of the central rest-mass density for the two schemes in the right panel of Figure 21. We find that the low-density fluctuations appearing in the stellar exterior with the EL5 scheme do not impact the solution in the stellar interior, with the low-frequency central density oscillations essentially being in phase for the two schemes. Also quite apparent is that the EL5 scheme yields rather constant-amplitude oscillations and this should be contrasted with the MP5 scheme, where the oscillations are comparatively larger in the first



$\sim 2,000M_{\odot}$ of the evolution. In both cases, however, the oscillations are extremely small and below 0.1%.

4.2.6 Grazing collision of neutron stars

We further test the ELH scheme in another truly dynamical test: the motion across the numerical grid of two neutron stars in a grazing collision. This is a setup that is very similar to that of a binary-neutron star system in quasi-circular orbit, the most obvious difference being the initial momenta of the two stars do not result in a quasi-circular orbit and that the initial fluid velocity can be taken to be arbitrary. In practice, the initial data is set up by generating two identical TOV models (the same as considered in Sections 4.2.1 and 4.2.2), superimposing the two data sets on the computational grid and imparting suitable initial momenta resulting in a small, but nonzero, impact parameter. Clearly, such initial data is valid only as a first approximation since the stars are not in the hydrostatic equilibrium corresponding to the binary system and the initial metric and extrinsic curvature do not reflect a solution of the Einstein constraints equations.

These violations lead to initial oscillations in the evolution (see Kastaun et al. 2013; Tsatsin and Marronetti 2013 for a more detailed discussion of a more sophisticated setup in which the stars are also subject to a spin up) which can however be reduced significantly by setting the initial distance of the two stars to a rather large value. More importantly, however, these oscillations do not interfere with the main goal of this test, namely, that of validating the ability of the ELH scheme to preserve sharply the features

of the stellar surface also when the star moves across the numerical grid.

More in detail, the star centers are set at positions $(x_1, y_1, z_1) = (50, -50, 0)$ and $(x_2, y_2, z_2) = (-50, 50, 0)$ in units of M_{\odot} , i.e., symmetric with respect to the grid center on the (x, y) plane and at a distance of $\sim 141 M_{\odot}$. The initial 3-velocities are $(v_1^x, v_1^y, v_1^z) = (0, -0.1, 0)$ and $(v_2^x, v_2^y, v_2^z) = (0, 0.1, 0)$ respectively. We evolve the system on a cubic grid of radius $512 M_{\odot}$, but employ reflection symmetry boundary conditions across the (x, y) plane and 180 degrees rotation symmetry boundary conditions across the (y, z) plane to reduce the computational cost. The grid structure consists of two identical box-in-box refinement levels hierarchies with refinement factor 2, each centered on a star and consisting of 5 cubic levels with radii 12, 25, 50, 100, 200 M_{\odot} , plus the coarse base level with radius $512 M_{\odot}$, so that the grid spacing in the innermost refinement level is $\Delta^i = 0.2 M_{\odot} \simeq 0.3$ km. The refinement levels moved to track the positions of the stars during the evolution (see also Radice et al. 2016 for further details on the initial data and grid structure). We set again $\Delta t = 0.15 \Delta x$.

The evolution consists in the two stars initially traversing the grid in the x direction and approaching each other, then bending their trajectories as in a gravitational scattering process; we do not follow the dynamics of the process after the first fly-by. Figure 22 shows the rest-mass density distribution on the (x, y) plane for this test. The snapshots of one of the two stars are taken at time $t = 768 M_{\odot}$, when the two stars are past the point of closest approach and are




Figure 22 Grazing collision test: density on (x, y) plane. Two-dimensional rest-mass density distribution relative to the initial data maximum value on the equatorial (x, y) plane for the grazing-collision tests at time $t = 768 M_{\odot}$, i.e., after the point of closest approach.

escaping. One can appreciate the deformation due to the boost, the acquired spin angular momentum and the tidal gravitational interaction. From the hydrodynamics point of view the behaviour of the MP5 and EL5 schemes is consistent with the previous tests, in particular the TOV in

a dynamical spacetime: EL5 shows a sharper star surface with respect to MP5, as well as an ‘emptier’ surrounding region, while the bulk of the star itself is very well resolved by both schemes.

The conservation of the total rest mass (left panel of Figure 23) is instead similar to the rotating-star case, thus showing a better performance by the MP5 scheme. Note however that the differences between the schemes are far smaller in this case, less than one order of magnitude. Furthermore, in contrast with the preceding tests, the grid structure in this case is much more complicated as well as dynamically updated to track the stars. Interpolation errors at the refinement level boundaries play therefore a greater role in the conservation of rest mass.

The evolution of the rest-mass density at the star centers, as shown in the right panel of Figure 23, is very similar for both schemes. There is an initial sudden increase in the density of about 4% with respect to the initial value, due to the evolution scheme bringing the star in hydrostatic equilibrium from the initial state. The density then oscillates around this new value, due to perturbations that in this case are not only induced by the violation of the constraint equations but also by the gravitational interaction. Overall, both schemes reproduce well all of these effects and show a very good agreement.

4.2.7 Gravitational collapse to a black hole

As a final test we evolve the violently dynamical collapse of a star to a black hole. This is also a common numerical-relativity benchmark (see, e.g., Font et al. 2002; Baiotti




Figure 23 Grazing collision test: total mass and central density evolution. Deviation of the total rest mass (left panel) and central rest-mass density (right panel) from the initial values for the grazing-collision test.

et al. 2005; Baiotti and Rezzolla 2006; Thierfelder et al. 2010), which allows us to validate ELH in the presence of a physical singularity and of an apparent horizon. More specifically, we consider a nonrotating star with central rest-mass density $8 \times 10^{-3} M_{\odot}^{-2}$, corresponding to a baryon mass of $1.5 M_{\odot}$ and radius $6 M_{\odot}$, and initiate the collapse with a velocity perturbation analogous to the one used in the migration test, but with the opposite sign, i.e., radially ingoing.

We define the time of black-hole formation as the instant at which an apparent horizon is first detected on the numerical domain, which, given the chosen setup, happens at $t \simeq 48 M_{\odot}$. Since we use singularity avoiding slicing conditions, we do not need to excise the interior spacetime of the black hole (Baiotti and Rezzolla 2006; Baiotti et al. 2007; Thierfelder et al. 2010). At the same time, we set the hydrodynamical variables to their atmosphere values inside a surface with the same shape as the apparent horizon, but radius $r = 0.9 r_{\text{AH}}$ in every angular direction, r_{AH} being the radius of the apparent horizon. This ‘hydrodynamic excision’ is not strictly necessary as our code can handle the collapse without it, regardless of the scheme we employ. However, we have observed that its use improves the accuracy of the subsequent evolution, most notably, it improves the behaviour of the rest-mass density and we therefore choose to employ it.

Figure 24 shows a snapshot of the rest-mass density on the equatorial plane just after the collapse. The central area of uniform low density is where the hydrodynamical vari-




Figure 24 BH collapse test: density on (x, y) plane.

Two-dimensional rest-mass density distribution relative to the initial data maximum value on the equatorial (x, y) plane at time $t = 61 M_{\odot}$ for the collapse test. Streams of matter ejected during collapse and accreting back onto the black hole are clearly visible and more prominent in EL5 data.

ables have been set to atmosphere inside the horizon, and in this plot appears of identical size and shape for the two codes. The areas outside the horizon are instead filled with matter which has been spuriously ejected from the outer layers of the star during the collapse. The rest-mass density is evidently higher in the case of the EL5 scheme, due to a slightly higher proportion of matter ejected during the collapse, which is in turn triggered by larger oscillations around the stellar surface for the EL5 scheme. However, in both cases the total rest mass outside the horizon is tiny, $\sim 10^{-6} M_{\odot}$ for the EL5 scheme and $\sim 10^{-9} M_{\odot}$ for MP5, and thus dynamically essentially irrelevant on the properties of the solution. Furthermore, most of this matter is gravitationally bound and hence accretes back onto the newly formed black hole, forming streams of infalling matter. This is particularly evident along the coordinate directions, where the numerical viscosity of the high-order finite-differences stencil is smaller, independently of the scheme employed.

The very close agreement between the two solutions is summarised in Figure 25, where the evolution of the central rest-mass density, minimum lapse and L_2 -norm of the Hamiltonian constraint is plotted. The discontinuities in the curves at about $50 M_{\odot}$ correspond to the time of collapse and arise because we exclude points inside the horizon in the calculation of extrema and norms. In each panel before the formation of the apparent horizon, the curves


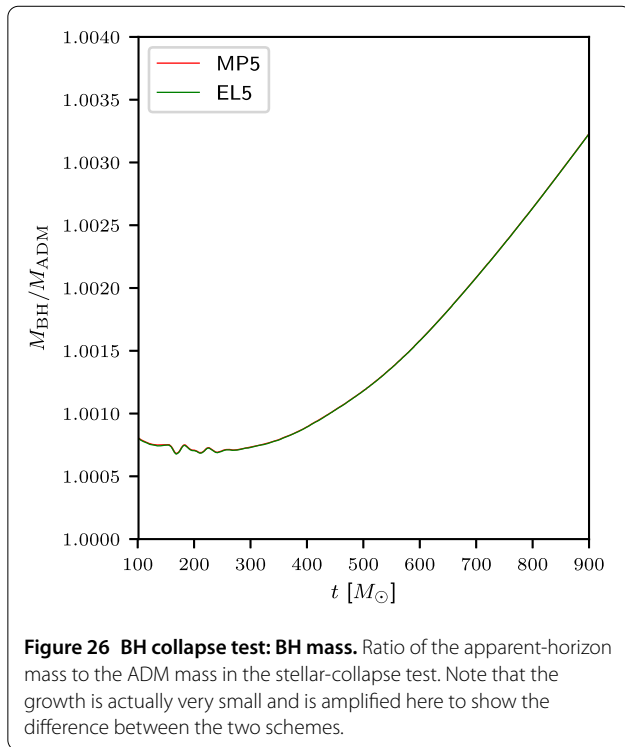


Figure 25 BH collapse test: density, lapse and Hamiltonian constraint evolution.

Maximum rest-mass density (top), minimum lapse (middle) and L_2 -norm of the Hamiltonian constraint (top) for the stellar-collapse test. Note that the violation of the Hamiltonian constraint will grow on longer timescales as it is typical of BSSNOK evolutions (Alic et al. 2013).



corresponding to the EL5 and MP5 schemes are essentially on top of each other (the largest differences being of the order of 0.3%, 0.6%, and 4.7% for each plot, respectively), showing the very good agreement in the evolution between the two schemes. Note also that after the apparent horizon formation and due to our approach of ‘hydrodynamic excision’, the upper panel of Figure 25 shows the maximum of the density in the exterior of the horizon rather than the central density. The disagreement in the EL5 and MP5 curves relates therefore to the tiny amount of residual matter outside of the black hole, and as such it has no dynamical impact on the black-hole solution.

A final confirmation of the equivalence between the two numerical solutions comes from the comparison of the two black-hole masses computed using the dynamical horizon formalism (Ashtekar and Krishnan 2003) and shown in Figure 26. As can be seen from the figure, we find a very close agreement between the two schemes.

5 Conclusions

We have presented a new high-order numerical method for the solution of the Euler equations of general-relativistic hydrodynamics that we name ‘entropy-limited hydrodynamics’ (ELH). The scheme is of the flux-limiting type, where a high-order numerical flux is combined with a stable low-order method, namely the Lax-Friedrichs flux. The flux-limiting is activated and driven by a shock indicator based on a measure of the entropy generated by the solution. Such a special and general-relativistic method is inspired by the entropy-viscosity method proposed recently

for the solution of the classical equations of hydrodynamics (Guermont et al. 2011), but it is also importantly different in that it does not require any change in the equations of relativistic hydrodynamics.

To assess the robustness and accuracy of our new method, we have discussed its implementation in the `WhiskyTHCEL` code, which exploits the finite-difference capabilities of the `WhiskyTHC` code, and tested its validity with an extensive series of tests, comparing the results of ELH with those obtained with another well-tested and high-order HRSC scheme: the fifth-order monotonicity-preserving MP5 method (Suresh and Huynh 1997). Overall, we have found that the scheme is stable and able to cope with shocks and discontinuities, both in classical test such as shock-tube tests, as well as in realistic astrophysical simulations.

Under all of these conditions the scheme has been found to be stable and to yield accuracy that is comparable, if not better, of that of the MP5 method. In some tests involving stars that are nonrotating or not moving across the grid, it also offers definite advantages, such as a sharper resolution of the surface/vacuum interface. At the same time, however, it also shows a less good conservation of the rest mass for stars that are rotating or moving across the computational domain (the opposite is true for stationary nonrotating stars, where the new method conserves rest mass more accurately). Quite surprisingly, all of the results presented here were obtained without any fine tuning of the two arbitrary coefficients that enter the definition of the scheme. Finally, thanks to its linearity and simplicity, the ELH method can also offer advantages in terms of performance. In our tests, we have found EL5 to be $\sim 50\%$ faster than MP5, even though our implementation was not particularly optimized. For instance, a definite advantage of ELH, which we did not exploit, is that it can be easily vectorised. At the same time, we remark that the exact speed-up that can be achieved with ELH depends also on external factors, such as the grid setup and number of ghost zones, which may vary for different applications. An interesting development in this sense would be the use of this scheme in a discontinuous Galerkin framework, whose superior scalability properties should decouple the performance of the ELH method from the grid setup.

The work presented here could be improved in at least two ways. Firstly, the already good capturing properties of steep gradients as those given by the stellar surface, could be further enhanced and the full capabilities of the scheme further exploited by coupling it to truly multidimensional stencils. Second, the two free coefficients that appear in the method, and that we have here set to unity for simplicity, could potentially be tuned to optimise some of the features of the solution. Both of these aspects will be explored in future work.

In conclusion, we have shown that entropy-limited hydrodynamics is a robust, stable, and accurate alternative

to commonly employed HRSC schemes. Its performance reaches the level of accuracy and stability necessary to apply it to realistic astrophysical simulations. Given these encouraging prospects, work is already in progress to apply this method to realistic simulations of binaries involving neutron stars and black holes.

Appendix 1: General relativistic hydrodynamics formulation

In this section we summarize the general relativistic formulation of hydrodynamics, which forms the theoretical background for the neutron star tests presented above; as well as providing details on the ELH scheme in the general relativistic case, when they differ from the discussion of Section 3.

A.1 3 + 1 decomposition

We adopt the usual 3 + 1 decomposition of spacetime (see Alcubierre 2008; Bona et al. 2009; Baumgarte and Shapiro 2010; Gourgoulhon 2012; Rezzolla and Zanotti 2013; Shibata 2016 for a detailed discussion), in which the spacetime is decomposed into spacelike hypersurfaces with normal $n^\mu = \alpha^{-1}(1, -\beta^i)$, where α is the lapse function and β^i are the components of the shift vector. Within this formalism, the spacetime metric $g_{\mu\nu}$ is written as

$$\begin{aligned} ds^2 &= g_{\mu\nu} dx^\mu dx^\nu \\ &= -\alpha^2 dt^2 + \gamma_{ij}(dx^i + \beta^i dt)(dx^j + \beta^j dt), \end{aligned} \quad (29)$$

where γ_{ij} is the spatial three-metric, which together with the extrinsic curvature $K_{ij} = -\frac{1}{2}\mathcal{L}_n\gamma_{ij}$, \mathcal{L}_n being the Lie derivative along n^μ , fully determines the geometry of each leaf of the foliation.

The matter content of the spacetime is described through its energy-momentum tensor $T_{\mu\nu}$, which within the 3 + 1 split can be decomposed in its timelike, spacelike and mixed parts as (see e.g., Rezzolla and Zanotti 2013)

$$E := n^\mu n^\nu T_{\mu\nu}, \quad (30)$$

$$S_i := -\gamma_i^\mu n^\nu T_{\mu\nu}, \quad (31)$$

$$S_{ij} := \gamma_i^\mu \gamma_j^\nu T_{\mu\nu}, \quad (32)$$

where $\gamma_\nu^\mu := \delta_\nu^\mu + n_\nu n^\mu$ is the projection operator orthogonal to n_μ , hence yielding purely spatial tensors.

A.2 Relativistic hydrodynamics

The general relativistic generalization of the energy-momentum tensor (1) consists in replacing the Minkowski metric $\eta_{\mu\nu}$ with the generic metric $g_{\mu\nu}$, i.e.

$$T_{\mu\nu} = \rho h u_\mu u_\nu + p g_{\mu\nu}. \quad (33)$$

The corresponding equations of motion for the fluid are the ‘conservation’ of the stress-energy tensor

$$\nabla_\mu T^{\mu\nu} = 0, \quad (34)$$

and conservation of rest mass

$$\nabla_\mu (\rho u^\mu) = 0, \quad (35)$$

which differ from their special relativistic counterparts (2) and (3) by the use of the ∇ operator, i.e. the covariant derivative constructed from the metric (29).

Equations (34) and (35) are solved in conservation form (Eq. (5)), where the ‘conserved’ variables \mathbf{U} are defined as

$$\begin{aligned} \mathbf{U} &:= \sqrt{\gamma} \begin{pmatrix} D \\ S_j \\ \tau \end{pmatrix} \\ &:= \sqrt{\gamma} \begin{pmatrix} \rho W \\ \rho h W^2 v_j \\ \rho h W^2 - p - \rho W \end{pmatrix}. \end{aligned} \quad (36)$$

This definitions differ from (6) by the multiplicative factor γ , i.e. the determinant of the 3-metric γ_{ij} . The fluxes and sources, containing metric-dependent terms, are given by

$$\mathbf{F}^i = \sqrt{\gamma} \begin{pmatrix} (\alpha v^i - \beta^i) D \\ \alpha S_j^i - \beta^i S_j \\ \alpha (S^i - D v^i) - \beta^i \tau \end{pmatrix}, \quad (37)$$

and

$$\mathbf{S} = \sqrt{\gamma} \begin{pmatrix} 0 \\ \frac{1}{2} \alpha S^{lm} \partial_j \gamma_{lm} + S_k \partial_j \beta^k - E \partial_j \alpha \\ \alpha S^{ij} K_{ij} - S^k \partial_k \alpha \end{pmatrix}. \quad (38)$$

This formulation is known as the ‘Valencia formulation’ and was first proposed by (Banyuls et al. 1997). Note that the sources terms for the momentum and energy equations are non-vanishing, which corresponds to the fact that the momentum and energy of the fluid are not independently conserved, but the coupling of the fluid to the spacetime and vice versa has to be taken into account (see e.g., Rezzolla and Zanotti 2013; Shibata 2016; Baumgarte and Shapiro 2010 for details). The fluid three-velocity measured by the normal observers is defined as

$$v^i := \frac{1}{\alpha} \left(\frac{u^i}{u^t} + \beta^i \right) \quad (39)$$

which also contains metric-dependent terms, and the Lorentz factor is $W := (1 - v^i v_i)^{-\frac{1}{2}} = \alpha u^t$. We also have used the fact that $\sqrt{-g} = \alpha \sqrt{\gamma}$.

A.3 ELH scheme extension to GR

The general relativistic formulation of the second law of thermodynamics, and therefore of the entropy residual \mathcal{R} , is

$$\mathcal{R} := \nabla_\mu (\rho s u^\mu) \geq 0, \quad (40)$$

once again swapping the partial derivative ∂ for the covariant one ∇ . This gets rewritten as

$$\begin{aligned} \mathcal{R} &= \nabla_\mu (s \rho u^\mu) \\ &= s \nabla_\mu (\rho u^\mu) + \rho u^\mu \nabla_\mu s \\ &= \rho u^\mu \nabla_\mu s = \rho u^\mu \partial_\mu s, \end{aligned} \quad (41)$$

which again involves only derivatives of the specific entropy s , and where the continuity equation (35) was used to remove the first term on the second line. The 4-velocity u^μ can again be written in terms of the fluid three-velocity v^i , but it will also contain the lapse and shift (see Eq. (39)), so that the final form of the residual in the 3 + 1 split is

$$\mathcal{R} = \frac{\rho W}{\alpha} [\partial_t s + (\alpha v^i - \beta^i) \partial_i s]. \quad (42)$$

Appendix 2: Spacetime evolution

The evolution of the spacetime is of course determined by Einstein's equations

$$R_{\mu\nu} - \frac{1}{2} R g_{\mu\nu} = 8\pi T_{\mu\nu}, \quad (43)$$

where $R_{\mu\nu}$ is the Ricci tensor and $R := R^\mu_\mu$ the Ricci scalar. The energy-momentum tensor $T_{\mu\nu}$ given by (33) appears in the right-hand-side of the equations, coupling the spacetime to the and fluid and vice versa.

We employ the common BSSNOK formulation (Shibata and Nakamura 1995; Baumgarte and Shapiro 1999; Brown 2009) of the Einstein equations (extensions to the CCZ4 (Alic et al. 2012) system are also in progress), along with the standard 1+log slicing condition and Gamma driver to specify the evolution of the lapse and shift, respectively (see Baiotti and Rezzolla 2016 for a review of these gauges). Given the following definitions for the evolved fields

$$\phi = \frac{1}{12} \ln(\gamma), \quad (44a)$$

$$K = \gamma^{ij} K_{ij}, \quad (44b)$$

$$\tilde{\gamma}_{ij} = e^{-4\phi} \gamma_{ij}, \quad (44c)$$

$$\tilde{A}_{ij} = e^{-4\phi} \left(K_{ij} - \frac{1}{3} \gamma_{ij} K \right), \quad (44d)$$

$$\tilde{\Gamma}^i = \tilde{\gamma}^{jk} \tilde{\Gamma}_{jk}^i, \quad (44e)$$

where $\tilde{\Gamma}_{jk}^i$ are the Christoffel symbols computed from the conformal metric $\tilde{\gamma}_{ij}$, the BSSNOK and gauge equations take the form:

$$\partial_\perp \phi = \frac{1}{6} \partial_k \beta^k - \frac{1}{6} \alpha K, \quad (45a)$$

$$\partial_\perp \tilde{\gamma}_{ij} = -2\alpha \tilde{A}_{ij} - \frac{2}{3} \tilde{\gamma}_{ij} \partial_k \beta^k, \quad (45b)$$

$$\begin{aligned} \partial_\perp K &= \alpha \left(\tilde{A}_{ij} \tilde{A}^{ij} + \frac{1}{3} K^2 \right) - \gamma^{ij} \nabla_i \nabla_j \alpha \\ &\quad + 4\pi (S_k^k + E), \end{aligned} \quad (45c)$$

$$\begin{aligned} \partial_\perp \tilde{A}_{ij} &= e^{-4\phi} [\alpha (R_{ij} - 8\pi S_{ij}) - \nabla_i \nabla_j \alpha]^{\text{TF}} \\ &\quad - \frac{2}{3} \tilde{A}_{ij} \partial_k \beta^k + \alpha (K \tilde{A}_{ij} - 2 \tilde{A}_{ik} \tilde{A}_j^k), \end{aligned} \quad (45d)$$

$$\begin{aligned} \partial_\perp \tilde{\Gamma}^i &= \tilde{\gamma}^{kl} \partial_k \partial_l \beta^i + \frac{2}{3} \tilde{\gamma}^{jk} \tilde{\Gamma}_{jk}^i \partial_l \beta^l \\ &\quad + \frac{1}{3} \tilde{\nabla}^i (\partial_k \beta^k) - 2 \tilde{A}^{ik} \partial_k \alpha + 2 \alpha \tilde{A}^{kl} \tilde{\Gamma}_{kl}^i \\ &\quad + 12 \alpha \tilde{A}^{ik} \partial_k \phi - \frac{4}{3} \alpha \tilde{\nabla}^i K - 16 \pi \alpha \tilde{\gamma}^{ij} S_j, \end{aligned} \quad (45e)$$

$$\partial_t \alpha = -2\alpha K + \beta^k \partial_k \alpha, \quad (45f)$$

$$\partial_t \beta^i = \frac{3}{4} B^i + \beta^k \partial_k \beta^i, \quad (45g)$$

$$\partial_t B^i = \partial_t \tilde{\Gamma}^i - \eta B^i + \beta^k \partial_k B^i, \quad (45h)$$

where the operator ∂_\perp stands for $\partial_t - \mathcal{L}_\beta$, i.e., the derivative with respect to coordinate time minus the Lie derivative along the shift, and the notation $[\dots]^{\text{TF}}$ indicates terms that are made trace free with respect to the conformal metric. The covariant derivatives ∇ and $\tilde{\nabla}$ are constructed from the physical and covariant three metric respectively, and $\eta \sim 1/M$ (M being the mass of the system).

The three dimensional Ricci tensor R_{ij} is split in two parts, $R_{ij} = \tilde{R}_{ij}^\phi + \tilde{R}_{ij}$, the first involving the conformal factor ϕ and the second the derivatives of the conformal metric $\tilde{\gamma}_{ij}$:

$$\begin{aligned} \tilde{R}_{ij}^\phi &= \phi^{-2} [\phi (\tilde{\nabla}_i \tilde{\nabla}_j \phi + \tilde{\gamma}_{ij} \tilde{\nabla}^k \tilde{\nabla}_k \phi) \\ &\quad - 2 \tilde{\gamma}_{ij} \tilde{\nabla}^k \phi \tilde{\nabla}_k \phi], \end{aligned} \quad (46a)$$

$$\begin{aligned} \tilde{R}_{ij} &= -\frac{1}{2} \tilde{\gamma}^{lm} \partial_l \partial_m \tilde{\gamma}_{ij} + \tilde{\gamma}_{k(i} \partial_j) \tilde{\Gamma}^k \\ &\quad + \tilde{\Gamma}^k \tilde{\Gamma}_{(ij)k} + \tilde{\gamma}^{lm} [2 \tilde{\Gamma}^k_{l(i} \tilde{\Gamma}_{j)km} + \tilde{\Gamma}_{im}^k \tilde{\Gamma}_{kj}^l]. \end{aligned} \quad (46b)$$

In integrating these equations a constrained approach is used, i.e., we enforce the constraints $\det \tilde{\gamma}_{ij} = 1$ and $\text{tr} \tilde{A}_{ij} = 0$ at every step.

The spacetime evolution is taken care of by the `McLachlan` code (Brown et al. 2009) to evolve the spacetime variables in the BSSNOK formulation. `McLachlan`

approximates the equations using standard central finite-difference operators with upwinding of the shift advection terms and Kreiss-Oliger dissipation (Kreiss and Oliger 1973) to ensure stability. It supports up to eighth-order operators, however since the major source of errors in our simulations is the hydrodynamical part, we restrict ourselves here to fourth-order accuracy for the spacetime.

Acknowledgements

We thank Kentaro Takami for providing the stellar oscillation eigenfrequencies, while Erik Schnetter, Ian Hinder, and Massimiliano Leoni for useful discussions. The simulations were performed on the SuperMUC cluster at the LRZ in Garching, on the LOEWE cluster in CSC in Frankfurt, on the HazelHen cluster at the HLRS in Stuttgart and on the Caltech compute cluster Zwicky.

Funding

This research is supported in part by the ERC synergy grant 'BlackHoleCam: Imaging the Event Horizon of Black Holes' (Grant No. 610058), by 'NewCompStar', COST Action MP1304, by the LOEWE-Program in the Helmholtz International Center (HIC) for FAIR, and by the European Union's Horizon 2020 Research and Innovation Programme (Grant 671698) (call FETHPC-1-2014, project ExaHyPE). FG is supported by HIC for FAIR and the graduate school HGS-HIRE. DR gratefully acknowledges support from the Schmidt Fellowship and the Max-Planck/Princeton Center (MPPC) for Plasma Physics (NSF PHY-1523261).

Competing interests

The authors declare that they have no competing interests.

Authors' contributions

The implementation of the EL method in `whiskyTHC` and all of the tests were performed by FG. The project was initiated by DR and LR, and was closely supervised by both.

Author details

¹Institut für Theoretische Physik, Goethe Universität, Max-von-Laue-Str. 1, Frankfurt am Main, 60438, Germany. ²Institute for Advanced Study, 1 Einstein Dr., Princeton, NJ 08540, USA. ³Department of Astrophysical Sciences, Princeton University, 4 Ivy Lane, Princeton, NJ 08544, USA. ⁴Frankfurt Institute for Advanced Studies, Ruth-Moufang-Str. 1, Frankfurt am Main, 60438, Germany.

Endnotes

- Note that the operators (9) and (10) return approximations of the function h defined by $f_i =: \int_{x_i-1/2}^{x_i+1/2} h(x') dx'$ at $x_{i+1/2}$. In this sense, they act on volume averages, the point-wise flux being the volume average of h . The values $h_{i\pm 1/2}$ should appear in (8) instead of $f_{i\pm 1/2}$. We have here simplified the notation, but a full discussion can be found in Radice and Rezzolla (2012).
- Further details can be found in Radice et al. (2014a). Here, θ is computed following the algorithm in Section 2.2 of Hu et al. (2013), but replacing the rest-mass density ρ with its conserved relativistic counterpart D .
- The use of a higher-order stencil in the EL approach, e.g., EL7, does not yield to improvements in the solution; the treatment of the low-density regions is far more delicate and the mass conservation is degraded.
- Of course, for both schemes the amount of rest-mass outside the star is minute, being only 10^{-7} of the initial rest-mass for the EL5 scheme and $\sim 10^{-5}$ for the MP5 scheme.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Received: 19 December 2016 Accepted: 20 June 2017

Published online: 04 July 2017

References

- Alcubierre, M: Introduction to 3 + 1 Numerical Relativity. Oxford University Press, Oxford (2008). doi:10.1093/acprof:oso/9780199205677.001.0001
- Alic, D, Bona-Casas, C, Bona, C, Rezzolla, L, Palenzuela, C: Conformal and covariant formulation of the Z4 system with constraint-violation damping. *Phys. Rev. D* **85**(6), 064040 (2012). doi:10.1103/PhysRevD.85.064040, arXiv:1106.2254
- Alic, D, Kastaun, W, Rezzolla, L: Constraint damping of the conformal and covariant formulation of the Z4 system in simulations of binary neutron stars. *Phys. Rev. D* **88**(6), 064049 (2013). doi:10.1103/PhysRevD.88.064049, arXiv:1307.7391
- Anile, AM: Relativistic Fluids and Magneto-Fluids. Cambridge University Press, Cambridge (1990)
- Ashtekar, A, Krishnan, B: Dynamical horizons and their properties. *Phys. Rev. D* **68**, 104030 (2003). arXiv:gr-qc/0308033
- Baiotti, L, Damour, T, Giacomazzo, B, Nagar, A, Rezzolla, L: Accurate numerical simulations of inspiralling binary neutron stars and their comparison with effective-one-body analytical models. *Phys. Rev. D* **84**(2), 024017 (2011). doi:10.1103/PhysRevD.84.024017, arXiv:1103.3874
- Baiotti, L, Hawke, I, Montero, P, Rezzolla, L: A new three-dimensional general-relativistic hydrodynamics code. In: Capuzzo-Dolcetta, R (ed.) *Computational Astrophysics in Italy: Methods and Tools*, vol. 1, p. 210. MSAIt, Trieste (2003)
- Baiotti, L, Hawke, I, Montero, PJ, Löffler, F, Rezzolla, L, Stergioulas, N, Font, JA, Seidel, E: Three-dimensional relativistic simulations of rotating neutron-star collapse to a Kerr black hole. *Phys. Rev. D* **71**(2), 024035 (2005). doi:10.1103/PhysRevD.71.024035, arXiv:gr-qc/0403029
- Baiotti, L, Hawke, I, Rezzolla, L: On the gravitational radiation from the collapse of neutron stars to rotating black holes. *Class. Quantum Gravity* **24**, 187-206 (2007). arXiv:gr-qc/0701043
- Baiotti, L, Hawke, I, Rezzolla, L, Schnetter, E: Gravitational-wave emission from rotating gravitational collapse in three dimensions. *Phys. Rev. Lett.* **94**, 131101 (2005). arXiv:gr-qc/0503016
- Baiotti, L, Rezzolla, L: Challenging the paradigm of singularity excision in gravitational collapse. *Phys. Rev. Lett.* **97**, 141101 (2006). arXiv:gr-qc/0608113
- Baiotti, L, Rezzolla, L: Binary neutron-star mergers: a review of Einstein's richest laboratory. arXiv:1607.03540 (2016)
- Banyuls, F, Font, JA, Ibáñez, JM, Martí, JM, Miralles, JA: Numerical 3 + 1 general-relativistic hydrodynamics: a local characteristic approach. *Astrophys. J.* **476**, 221-231 (1997). doi:10.1086/303604
- Baumgarte, TW, Shapiro, SL: Numerical integration of Einstein's field equations. *Phys. Rev. D* **59**(2), 024007 (1999). doi:10.1103/PhysRevD.59.024007, arXiv:gr-qc/9810065
- Baumgarte, TW, Shapiro, SL: *Numerical Relativity: Solving Einstein's Equations on the Computer*. Cambridge University Press, Cambridge (2010). doi:10.1017/cbo9781139193344
- Bernuzzi, S, Dietrich, T: Gravitational waveforms from binary neutron star mergers with high-order WENO schemes in numerical relativity. arXiv:1604.07999 (2016)
- Bernuzzi, S, Nagar, A, Thierfelder, M, Brüggmann, B: Tidal effects in binary neutron star coalescence. *Phys. Rev. D* **86**(4), 044030 (2012). doi:10.1103/PhysRevD.86.044030, arXiv:1205.3403
- Bona, C, Palenzuela-Luque, C, Bona-Casas, C: *Elements of Numerical Relativity and Relativistic Hydrodynamics: From Einstein's Equations to Astrophysical Simulations*. Lecture Notes in Physics. Springer, Berlin (2009). <http://books.google.co.uk/books?id=KgPGHaCUaAYC>
- Brown, D, Diener, P, Sarbach, O, Schnetter, E, Tiglio, M: Turduckening black holes: an analytical and computational study. *Phys. Rev. D* **79**, 044023 (2009). doi:10.1103/PhysRevD.79.044023, arXiv:0809.3533 [gr-qc]
- Brown, DJ: Covariant formulations of Baumgarte, Shapiro, Shibata, and Nakamura and the standard gauge. *Phys. Rev. D* **79**(10), 104029 (2009). doi:10.1103/PhysRevD.79.104029
- Bugner, M, Dietrich, T, Bernuzzi, S, Weyhausen, A, Brüggmann, B: Solving 3D relativistic hydrodynamical problems with weighted essentially nonoscillatory discontinuous Galerkin methods. *Phys. Rev. D* **94**(8), 084004 (2016). doi:10.1103/PhysRevD.94.084004, arXiv:1508.07147
- Colella, P, Woodward, PR: The piecewise parabolic method (PPM) for gas-dynamical simulations. *J. Comput. Phys.* **54**(1), 174-201 (1984). doi:10.1016/0021-9991(84)90143-8
- Cordero-Carrión, I, Cerdá-Durán, P, Dimmelmeier, H, Jaramillo, JL, Novak, J,ourgoulhon, E: Improved constrained scheme for the Einstein equations:

- an approach to the uniqueness issue. *Phys. Rev. D* **79**(2), 024017 (2009). doi:10.1103/PhysRevD.79.024017, arXiv:0809.2325
- De Pietri, R, Feo, A, Maione, F, Löffler, F: Modeling equal and unequal mass binary neutron star mergers using public codes. *Phys. Rev. D* **93**(6), 064047 (2016). doi:10.1103/PhysRevD.93.064047, arXiv:1509.08804
- DeBuhr, J, Zhang, B, Anderson, M, Neilsen, D, Hirschmann, EW: Relativistic hydrodynamics with wavelets. arXiv:1512.00386 (2015)
- Einstein Toolkit: open software for relativistic astrophysics (2010). <http://einstein toolkit.org>
- Font, JA: Numerical hydrodynamics and magnetohydrodynamics in general relativity. *Living Rev. Relativ.* **11**, 7 (2008). doi:10.12942/lrr-2008-7
- Font, JA, Goodale, T, Iyer, S, Miller, M, Rezzolla, L, Seidel, E, Stergioulas, N, Suen, W-M, Tobias, M: Three-dimensional numerical general relativistic hydrodynamics. II. Long-term dynamics of single relativistic stars. *Phys. Rev. D* **65**(8), 084024 (2002). doi:10.1103/PhysRevD.65.084024, arXiv:gr-qc/0110047
- Gottlieb, S, Ketcheson, D, Shu, C-W: High order strong stability preserving time discretizations. *J. Sci. Comput.* **38**, 251-289 (2009). doi:10.1007/s10915-008-9239-z
- Gourgoulhon, E: 3 + 1 Formalism in General Relativity. Lecture Notes in Physics, vol. 846. Springer, Berlin (2012). doi:10.1007/978-3-642-24525-1
- Guermont, JL, Pasquetti, R: Entropy-based nonlinear viscosity for Fourier approximations of conservation laws. *C. R. Math. Acad. Sci.* **346**, 801-806 (2008). doi:10.1016/j.crma.2008.05.013
- Guermont, JL, Pasquetti, R, Popov, B: Entropy viscosity method for nonlinear conservation laws. *J. Comput. Phys.* **230**(11), 4248-4267 (2011). doi:10.1016/j.jcp.2010.11.043
- Harten, A, Engquist, B, Osher, S, Chakravarthy, SR: Uniformly high order accurate essentially non-oscillatory schemes III. *J. Comput. Phys.* **71**, 231-303 (1987). doi:10.1016/0021-9991(87)90031-3
- Hu, XY, Adams, NA, Shu, C-W: Positivity-preserving method for high-order conservative schemes solving compressible Euler equations. *J. Comput. Phys.* **242**, 169-180 (2013). doi:10.1016/j.jcp.2013.01.024, arXiv:1203.1540
- Jiang, G-S, Shu, C-W: Efficient implementation of weighted eno schemes. *J. Comput. Phys.* **126**, 202-228 (1996)
- Kastaun, W, Galeazzi, F, Alic, D, Rezzolla, L, Font, JA: Black hole from merging binary neutron stars: how fast can it spin? *Phys. Rev. D* **88**(2), 021501 (2013). doi:10.1103/PhysRevD.88.021501, arXiv:1301.7348
- Kidder, LE, Field, SE, Foucart, F, Schnetter, E, Teukolsky, SA, Bohn, A, Deppe, N, Diener, P, Hébert, F, Lippuner, J, Miller, J, Ott, CD, Scheel, MA, Vincent, T: SpECTRE: A task-based discontinuous Galerkin code for relativistic astrophysics. arXiv:1609.00098 (2016)
- Kreiss, HO, Olinger, J: Methods for the Approximate Solution of Time Dependent Problems. GARP Publication Series, vol. 10, Geneva (1973)
- Lele, SK: Compact finite difference schemes with spectral-like resolution. *J. Comput. Phys.* **103**, 16-42 (1992)
- Leveque, RJ: Numerical Methods for Conservation Laws. Birkhäuser, Basel (1992)
- Liu, X-D, Osher, S, Chan, T: Weighted essentially non-oscillatory schemes. *J. Comput. Phys.* **115**, 200-212 (1994). doi:10.1006/jcph.1994.1187
- Löffler, F, Faber, J, Bentivegna, E, Bode, T, Diener, P, Haas, R, Hinder, I, Mundim, BC, Ott, CD, Schnetter, E, Allen, G, Campanelli, M, Laguna, P: The Einstein Toolkit: a community computational infrastructure for relativistic astrophysics. *Class. Quantum Gravity* **29**(11), 115001 (2012). doi:10.1088/0264-9381/29/11/115001, arXiv:1111.3344
- Martí, JM, Müller, E: Extension of the piecewise parabolic method to one-dimensional relativistic hydrodynamics. *J. Comput. Phys.* **123**, 1-14 (1996). doi:10.1006/jcph.1996.0001
- Martí, JM, Müller, E: Numerical hydrodynamics in special relativity. *Living Rev. Relativ.* **6**, 7 (2003). doi:10.12942/lrr-2003-7
- Martí, JM, Müller, E: Grid-based methods in relativistic hydrodynamics and magnetohydrodynamics. *Living Rev. Comput. Astrophys.* **1**, 3 (2015). doi:10.1007/lrca-2015-3
- Mignone, A, Tzeferacos, P, Bodo, G: High-order conservative finite difference GLM-MHD schemes for cell-centered MHD. *J. Comput. Phys.* **229**, 5896-5920 (2010). doi:10.1016/j.jcp.2010.04.013, arXiv:1001.2832
- Miller, JM, Schnetter, E: An operator-based local discontinuous Galerkin method compatible with the BSSN formulation of the Einstein equations. arXiv:1604.00075 (2016)
- Paschalidis, V: General relativistic simulations of compact binary mergers as engines of short gamma-ray bursts. arXiv:1611.01519 (2016)
- Radice, D, Galeazzi, F, Lippuner, J, Roberts, LF, Ott, CD, Rezzolla, L: Dynamical mass ejection from binary neutron star mergers. *Mon. Not. R. Astron. Soc.* **460**, 3255-3271 (2016). doi:10.1093/mnras/stw1227, arXiv:1601.02426
- Radice, D, Rezzolla, L: Discontinuous Galerkin methods for general-relativistic hydrodynamics: formulation and application to spherically symmetric spacetimes. *Phys. Rev. D* **84**(2), 024010 (2011). doi:10.1103/PhysRevD.84.024010, arXiv:1103.2426
- Radice, D, Rezzolla, L: THC: a new high-order finite-difference high-resolution shock-capturing code for special-relativistic hydrodynamics. *Astron. Astrophys.* **547**, 26 (2012). doi:10.1051/0004-6361/201219735, arXiv:1206.6502
- Radice, D, Rezzolla, L, Galeazzi, F: High-order fully general-relativistic hydrodynamics: new approaches and tests. *Class. Quantum Gravity* **31**(7), 075012 (2014a). doi:10.1088/0264-9381/31/7/075012, arXiv:1312.5004
- Radice, D, Rezzolla, L, Galeazzi, F: Beyond second-order convergence in simulations of binary neutron stars in full general-relativity. *Mon. Not. R. Astron. Soc. Lett.* **437**, 46-50 (2014b). doi:10.1093/mnras/slt137, arXiv:1306.6052
- Radice, D, Rezzolla, L, Galeazzi, F: High-order numerical-relativity simulations of binary neutron stars. In: Numerical Modeling of Space Plasma Flows: ASTRONUM-2014. ASP Conference Series, vol. 498, pp. 121-126 (2015). arXiv:1502.00551
- Radice, D, Rezzolla, L, Kellerman, T: Critical phenomena in neutron stars: I. Linearly unstable nonrotating models. *Class. Quantum Gravity* **27**(23), 235015 (2010). doi:10.1088/0264-9381/27/23/235015, arXiv:1007.2809.
- Read, JS, Baiotti, L, Creighton, JDE, Friedman, JL, Giacomazzo, B, Kyutoku, K, Markakis, C, Rezzolla, L, Shibata, M, Taniguchi, K: Matter effects on binary neutron star waveforms. *Phys. Rev. D* **88**(4), 044042 (2013). doi:10.1103/PhysRevD.88.044042, arXiv:1306.4065
- Rezzolla, L, Zanotti, O: Relativistic Hydrodynamics. Oxford University Press, Oxford (2013). doi:10.1093/acprof:oso/9780198528906.001.0001
- Roe, PL: Some contributions to the modelling of discontinuous flows. In: Lee, RL, Sani, RL, Shih, TM, Gresho, PM (eds.) Large-Scale Computations in Fluid Mechanics, pp. 163-193 (1985)
- Sagaut, P, Grohens, R: Discrete filters for large eddy simulation. *Int. J. Numer. Methods Fluids* **31**, 1195-1220 (1999). doi:10.1103/PhysRevD.79.104029
- Shibata, M: Numerical Relativity. World Scientific, Singapore (2016). doi:10.1142/9692
- Shibata, M, Nakamura, T: Evolution of three-dimensional gravitational waves: harmonic slicing case. *Phys. Rev. D* **52**, 5428-5444 (1995). doi:10.1103/PhysRevD.52.5428
- Shibata, M, Taniguchi, K: Coalescence of black hole-neutron star binaries. *Living Rev. Relativ.* **14**, 6 (2011). doi:10.12942/lrr-2011-6
- Shu, CW: Essentially non-oscillatory and weighted essentially non-oscillatory schemes for hyperbolic conservation laws. Lecture notes ICASE report 97-65; NASA CR-97-206253, NASA Langley Research Center. http://ntrs.nasa.gov/archive/nasa/casi.ntrs.nasa.gov/19980007543_1998045663.pdf (1997)
- Sod, GA: A survey of several finite difference methods for systems of nonlinear hyperbolic conservation laws. *J. Comput. Phys.* **27**, 1-31 (1978). doi:10.1016/0021-9991(78)90023-2
- Stergioulas, N, Friedman, JL: Comparing models of rapidly rotating relativistic stars constructed by two numerical methods. *Astrophys. J.* **444**, 306-311 (1995). doi:10.1086/175605, arXiv:astro-ph/9411032
- Suresh, A, Huynh, HT: Accurate monotonicity-preserving schemes with Runge-Kutta time stepping. *J. Comput. Phys.* **136**(1), 83-99 (1997). doi:10.1006/jcph.1997.5745
- Takami, K, Rezzolla, L, Yoshida, S: A quasi-radial stability criterion for rotating relativistic stars. *Mon. Not. R. Astron. Soc.* **416**, 1-5 (2011). doi:10.1111/j.1745-3933.2011.01085.x, arXiv:1105.3069
- Tchekhovskoy, A, McKinney, JC, Narayan, R: WHAM: a WENO-based general relativistic numerical scheme? I. Hydrodynamics. *Mon. Not. R. Astron. Soc.* **379**(2), 469-497 (2007). doi:10.1111/j.1365-2966.2007.11876.x
- Thierfelder, M, Bernuzzi, S, Brüggmann, B: Numerical relativity simulations of binary neutron stars. *Phys. Rev. D* **84**(4), 044012 (2011). doi:10.1103/PhysRevD.84.044012, arXiv:1104.4751
- Thierfelder, M, Bernuzzi, S, Hilditch, D, Brüggmann, B, Rezzolla, L: The trumpet solution from spherical gravitational collapse with puncture gauges. *Phys. Rev. D* **83**, 064022 (2010). arXiv:1012.3703
- Tsatsin, P, Marronetti, P: Initial data for neutron star binaries with arbitrary spins. *Phys. Rev. D* **88**, 064060 (2013). doi:10.1103/PhysRevD.88.064060

- Vichnevetsky, R, Bowles, JB: Fourier Analysis of Numerical Approximations of Hyperbolic Equations. SIAM, Philadelphia (1982)
- Wu, K: Design of provably physical-constraint-preserving methods for general relativistic hydrodynamics. *Phys. Rev. D* **95**, 103001 (2017). arXiv:1610.06274
- Yoshida, S, Eriguchi, Y: Quasi-radial modes of rotating stars in general relativity. *Mon. Not. R. Astron. Soc.* **322**, 389-396 (2001)
- Zanotti, O, Fambri, F, Dumbser, M: Solving the relativistic magnetohydrodynamics equations with ADER discontinuous Galerkin methods, a posteriori subcell limiting and adaptive mesh refinement. *Mon. Not. R. Astron. Soc.* **452**, 3010-3029 (2015). doi:10.1093/mnras/stv1510, arXiv:1504.07458
- Zhang, W, MacFadyen, AI: RAM: a relativistic adaptive mesh refinement hydrodynamics code. *Astrophys. J. Suppl. Ser.* **164**, 255-279 (2006)
- Zilhão, M, Löffler, F: An introduction to the Einstein toolkit. *Int. J. Mod. Phys. A* **28**, 40014 (2013). doi:10.1142/S0217751X13400149, arXiv:1305.5299
- Zingan, V, Guermond, JL, Morel, J, Popov, B: Implementation of the entropy viscosity method with the discontinuous Galerkin method. *Comput. Methods Appl. Mech. Eng.* **253**, 479-490 (2013). doi:10.1016/j.cma.2012.08.018

Submit your manuscript to a SpringerOpen[®] journal and benefit from:

- Convenient online submission
- Rigorous peer review
- Open access: articles freely available online
- High visibility within the field
- Retaining the copyright to your article

Submit your next manuscript at ► springeropen.com
