

Robust recovery of primitive variables in relativistic ideal magnetohydrodynamics

Wolfgang Kastaun^{1,2}, Jay Vijay Kalinani^{3,4}, and Riccardo Ciolfi^{5,4}

¹Max Planck Institute for Gravitational Physics (Albert Einstein Institute),
Callinstrasse 38, 30167 Hannover, Germany

²Leibniz Universität Hannover, 30167 Hannover, Germany

³Università di Padova, Dipartimento di Fisica e Astronomia,
Via Francesco Marzolo 8, I-35131 Padova, Italy

⁴INFN, Sezione di Padova, Via Francesco Marzolo 8, I-35131 Padova, Italy

⁵INAF, Osservatorio Astronomico di Padova, Vicolo dell'Osservatorio 5, I-35122 Padova, Italy



(Received 4 May 2020; revised 14 August 2020; accepted 2 October 2020; published 20 January 2021)

Modern simulation codes for general relativistic ideal magnetohydrodynamics are all facing a long-standing technical problem given by the need to recover fundamental variables from those variables that are evolved in time. In the relativistic case, this requires the numerical solution of a system of nonlinear equations. Although several approaches are available, none has proven completely reliable. A recent study comparing different methods showed that all can fail, in particular for the important case of strong magnetization and moderate Lorentz factors. Here, we propose a new robust, efficient, and accurate solution scheme, along with a proof for the existence and uniqueness of a solution, and analytic bounds for the accuracy. Further, the scheme allows us to reliably detect evolution errors leading to unphysical states and automatically applies corrections for typical harmless cases. A reference implementation of the method is made publicly available as a software library. The aim of this library is to improve the reliability of binary neutron star merger simulations, in particular in the investigation of jet formation and magnetically driven winds.

DOI: [10.1103/PhysRevD.103.023018](https://doi.org/10.1103/PhysRevD.103.023018)

I. INTRODUCTION

General relativistic magnetohydrodynamic (GRMHD) simulations are an important tool to study many astrophysical scenarios involving neutron stars (NSs) and black holes (BHs). In particular, they represent the leading approach to investigate the dynamics of binary neutron star (BNS) and NS-BH mergers, which are the most important events in the nascent field of multimessenger astrophysics with gravitational wave (GW) sources [1].

Arguably one of the most pressing unsolved problems related to BNS and NS-BH mergers is to find the exact mechanism powering short gamma-ray bursts (SGRBs). The simultaneous detection of the gravitational wave event GW170817 and the SGRB named GRB 170817A [2–4], along with the following “afterglow” emission across the entire electromagnetic spectrum (e.g., Refs. [1,5–7]),

provided compelling evidence that BNS mergers can power SGRBs (e.g., Refs. [8–10]). In turn, this implies that the remnant object resulting from a BNS merger can launch, at least in some cases, a relativistic jet, which was indeed confirmed for GRB 170817A [9,10]. However, the jet launching mechanism and the nature of the object acting as a central engine, either an accreting BH or a massive NS, remain uncertain (e.g., Ref. [11]).

Current simulations suggest that a mechanism based on neutrino-antineutrino annihilation would not be powerful enough to explain SGRBs [12,13], reinforcing the alternative idea that the main driver of jet formation should be a strong magnetic field. GRMHD simulations of BNS and NS-BH mergers, while considerably more complex and expensive because of the inclusion of magnetic fields, become necessary to properly address the problem. Recent studies in this direction already provide important hints, supporting a scenario where the central engine is an accreting BH [14,15] while disfavoring the massive NS scenario [16]. Nonetheless, a final answer is still missing, and it will be necessary to overcome the technical limitations of present GRMHD codes to ultimately solve the problem.

The merger event GW170817 was also accompanied by the kilonova transient AT 2017gfo, powered by the

Published by the American Physical Society under the terms of the *Creative Commons Attribution 4.0 International* license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Open access publication funded by the Max Planck Society.

radioactive decay of heavy r -process elements synthesized within the matter ejected by the merger (e.g., Refs. [1,17]). Although this kilonova was observed in unprecedented detail, the interpretation in terms of specific ejecta components and their physical origin is still under debate. Also in this case, numerical relativity simulations represent the ideal approach to fully understand the different mass ejection processes occurring in a BNS (or a NS-BH) merger. Moreover, for some of these ejection processes, magnetic fields are likely to play an important role (e.g., Refs. [18,19]), and therefore simulations should be performed in GRMHD.

The present work is devoted to a technical but crucial aspect of these simulations that has proven surprisingly difficult, and is motivated by the importance of GRMHD simulations in the context of BNS mergers (see, e.g., Ref. [20] for a recent review). Modern evolution codes are based on evolution equations written in the form of coupled conservation laws for baryon number density, energy, and momentum density, including the electromagnetic contributions, and either magnetic field or vector potential. Primitive variables such as matter velocity, density, and pressure, are not directly evolved. Instead, they have to be recovered from the evolved quasiconserved quantities after each evolution step.

While in Newtonian physics the above recovery is trivial, for the relativistic case one has to numerically solve a system of coupled nonlinear equations. The system involves also the equation of state (EOS), which encodes the behavior of matter up to supranuclear densities by specifying the pressure as a function of density and temperature. An additional degree of freedom is the electron fraction, which effectively describes the matter composition, and which can only change due to neutrino reactions. Since the EOS is not well constrained theoretically or by observation, a crucial requirement is the ability to perform simulations employing arbitrary EOSs. This precludes closed-form solutions for the primitive variables, and the system has to be solved numerically. Since the solution is required inside the innermost loop of the evolution, computational efficiency is almost as important as robustness.

Note that most evolution codes make the simplifying assumption of ideal MHD. Although the electrical conductivity in merger remnants is very high, this approximation might not be justified for all aspects of the problem. On the other hand, evolving resistive GRMHD equations introduces even more difficulties (see also Ref. [21]). The equations for the primitive variable recovery are also very different for resistive GRMHD. Another complication is that in regions with strong magnetic fields but low mass density, movement of the matter becomes dominated by the field. Treating this “force-free” regime would in principle require different numerical evolution methods (for example, see Ref. [22]).

The simpler problem of recovering the primitive variables in relativistic hydrodynamics without magnetic fields is already solved in a robust manner, as described in Ref. [23] (also adopted in Ref. [24]). For the full problem of ideal GRMHD, several recovery methods with different limitations have been employed in GRMHD evolution codes [25–32]. Older schemes such as Ref. [30] are limited to particular analytic prescriptions for the EOS. Newer schemes can in principle work with any EOS, but not all implementations actually allow the use of arbitrary EOSs. For a detailed review, we refer to Ref. [33].

All of the schemes investigated in Ref. [33] were shown to fail in certain regimes. While some of them work well enough in most of the regimes encountered during a merger simulation, even rare primitive recovery failures need to be handled and remain a common hurdle. An additional complication is that not all combinations of values for the evolved variables correspond to physically valid primitive variables. The occurrence of invalid evolved variables due to numerical errors of the evolution needs to be monitored and, if possible, corrected. If the recovery can fail also for valid input, it becomes impossible to reliably assess the overall validity.

In this work, we develop a new recovery algorithm with the mathematically proven ability to always find a solution, and which is guaranteed to recognize invalid evolved variables. Furthermore, the scheme provides mathematically derived accuracy bounds. Our scheme is limited to the ideal MHD approximation, but it does not introduce problems in the force-free regime. We provide a reference implementation which is ready to use in any GRMHD evolution code, in the form of a C++ library named `RePrimAnd` [34]. Our implementation is written to be completely EOS-agnostic and provides a framework for EOSs that can easily be extended. Since our aim is to improve reliability, we subject the numerical implementation of the algorithm to a comprehensive suite of tests, also studying the effects of finite floating-point precision.

The article is organized as follows: In Sec. II, we state the problem, derive the new scheme, prove the existence and uniqueness of a solution, and investigate the expected accuracy. In Sec. III, we discuss possible corrections to invalid evolved variables. In Sec. IV, we present numerical tests of our reference implementation, demonstrating robustness, efficiency, and precision. Here we also compare to other existing schemes. Then, we study error propagation of evolution errors to the primitive variables in Sec. V, identifying potentially problematic regimes. Finally, in Sec. VI we draw our conclusions.

II. FORMULATION OF THE SCHEME

A. Primitive variables

Our scheme is designed for evolution codes which evolve variables defined on a spacelike foliation of

spacetime from one time slice to the next. The hypersurfaces and their normal observers define a frame we will refer to as the Eulerian frame.

We denote the 3-velocity of the fluid with respect to the Eulerian frame as v^i , and the corresponding Lorentz factor as W . We will also use a quantity $z \equiv Wv$. The baryon number density in the fluid rest frame is denoted as n_B . It is common to multiply n_B with an arbitrary mass constant m_B to define the baryonic mass density $\rho = n_B m_B$. The pressure in the fluid rest frame is assumed to be isotropic and denoted as P . Denoting the fluid contribution to the total energy density in the fluid rest frame as ρ_E , we define the specific internal energy:

$$\epsilon = \frac{\rho_E}{\rho} - 1. \quad (1)$$

We further define $a = P/\rho_E$ and the relativistic enthalpy:

$$h = 1 + \epsilon + \frac{P}{\rho} = (1 + \epsilon)(1 + a). \quad (2)$$

Note that the definitions of ϵ and h both depend on the arbitrary choice of m_B .

The primitive variables we use to describe the electromagnetic field are the electric and magnetic fields as seen by an Eulerian observer. In terms of the Maxwell tensor,

$$E^\mu = n_\nu F^{\mu\nu}, \quad B^\mu = n_\nu {}^*F^{\mu\nu}, \quad (3)$$

where n is the normal to the hypersurfaces of the foliation, and the star denotes the Hodge dual. E^μ, B^μ are tangential to the hypersurface and thus equivalent to 3-vectors E^i, B^i . Our scheme neither requires nor provides the fields in the fluid frame, which can be obtained from the above using standard transformations.

B. Equation of state

We assume an equation of state (EOS) of the form

$$P = P(\rho, \epsilon). \quad (4)$$

The EOS could also depend on further variables, such as the electron fraction, as long as those variables are evolved variables or can be obtained from evolved variables in a trivial way, and can therefore be treated as fixed parameters in the primitive recovery algorithm.

For our purpose, it is also important to specify a validity range for each EOS. The validity range considers both physical and technical constraints. The most important physical constraint is the zero-temperature limit for the internal energy. An example of a technical constraint is the range of values available for an EOS given in tabulated form. Currently, our scheme uses an EOS-dependent validity region specified in the following form:

$$\rho_{\min} \leq \rho \leq \rho_{\max}, \quad (5)$$

$$\epsilon_{\min}(\rho) \leq \epsilon \leq \epsilon_{\max}(\rho). \quad (6)$$

However, it could easily be adapted to a more general shape in ρ, ϵ parameter space. We require that the lower validity bound $\epsilon_{\min}(\rho)$ be the zero-temperature value at the given density. This is the only meaningful choice for any numerical simulation involving cold matter at any time. Our error policy for correcting invalid evolved variables is based on this assumption, as is the proof for guaranteed success of the algorithm.

Our scheme relies on some physical constraints. Causality and thermodynamic stability require

$$0 \leq c_s^2 < 1, \quad (7)$$

where c_s is the adiabatic speed of sound, given by

$$c_s^2 = \left. \frac{d \ln(h)}{d \ln(\rho)} \right|_{s=\text{const}}. \quad (8)$$

If the EOS depends on electron fraction, it is also assumed to be constant in the above expression. We assume positive baryon number density and positive total energy density, which implies

$$0 \leq \rho_{\min} \leq \rho, \quad -1 < \epsilon_{\min}(\rho) \leq \epsilon. \quad (9)$$

We assume that the pressure is positive and further bounded by the total energy density (dominant energy condition), which implies

$$0 \leq a \leq 1. \quad (10)$$

For a given EOS, we also require a positive lower bound h_0 for the relativistic enthalpy h , such that $0 < h_0 \leq h(\rho, \epsilon)$ over the entire validity region of the EOS. This requirement only excludes exotic matter with $P \leq -\rho_E$.

Note that we do not assume $\epsilon > 0$ or $h \geq 1$. The definitions of ϵ and h depend on the arbitrary choice of the mass constant m_B . Unless the latter is fine-tuned to the average baryon binding energy at low density, nuclear physics EOSs often yield slightly negative ϵ . In practice, h_0 is of order unity.

By design, our scheme is not confined to any particular equation of state. It only uses the information defined above and does not make any other kind of EOS-specific distinctions or adjustments. For the purpose of testing our scheme, we use two specific EOSs as examples:

- (1) A hybrid EOS given by a cold part and a simple thermal part:

$$P(\rho, \epsilon) = P_{\text{cold}}(\rho) + (\Gamma_{\text{th}} - 1)\rho(\epsilon - \epsilon_{\text{cold}}(\rho)), \quad (11)$$

$$\epsilon_{\min}(\rho) = \epsilon_{\text{cold}}(\rho). \quad (12)$$

For the cold part, we employ the MS1 EOS from Ref. [35] (based on Ref. [36]), and we use $\Gamma_{\text{th}} = 1.8$. This hybridized type of EOS is often used in numerical relativity.

(2) The classical ideal gas, given by

$$P(\rho, \epsilon) = \rho\epsilon(\Gamma - 1), \quad (13)$$

$$\epsilon_{\min}(\rho) = 0. \quad (14)$$

Here, we use an adiabatic exponent $\Gamma = 2$. Pressure and internal energy are zero at zero temperature. We use this unrealistic model, where pressure is only given by thermal effects and degeneracy pressure is ignored, as a corner case for testing our algorithm. Note that in numerical relativity, an ideal gas EOS normally refers to a hybrid EOS based on a polytrope with adiabatic exponent $\Gamma = \Gamma_{\text{th}}$ as the zero-temperature limit.

We note that numerical relativity tools are often judged by how well they function with tabulated nuclear physics EOSs, because such tables often contain isolated points with thermodynamically inconsistent jumps or regimes with superluminal sound speed. Violations of basic physical constraints such as causality can lead to many fundamental problems with the evolution equations. In Sec. II G, we will point out a potential problem for primitive recovery. In our opinion, any effort to ensure that primitive variable recovery and other aspects of simulations can produce results with faulty EOSs is a step in the wrong direction. Instead, we advocate in favor of repairing minor EOS faults before use.

C. Evolved variables

Our scheme is intended for numerical evolution codes employing evolution equations for energy, momentum, and baryon number density formulated as a quasiconservation law. This is the case for finite-volume shock-capturing schemes. The evolved quantities are called conserved variables, although only the baryon number is strictly conserved. The fluid contribution is given by

$$\bar{D} = \rho W, \quad (15)$$

$$\bar{\tau} = \bar{D}(hW - 1) - P, \quad (16)$$

$$\bar{S}_i = \bar{D}W h v_i. \quad (17)$$

Including also EM contributions, the evolved variables are given by

$$D = \bar{D}, \quad (18)$$

$$\tau = \bar{\tau} + \frac{1}{2}(E^2 + B^2), \quad (19)$$

$$S_i = \bar{S}_i + \epsilon_{ijk} E^j B^k. \quad (20)$$

In most formulations, the evolved variables above include the volume element of the spatial metric. Since this factor is available from the spacetime evolution, it is not relevant for the following and was left out of the definitions.

In addition to the evolved variables D , τ , S_i , we assume that the magnetic field B in the Eulerian frame is either an evolved variable or can be reconstructed from evolved variables without knowledge of the fluid-related primitive variables, such that B is known when our primitive reconstruction scheme is run. This is the case for state-of-the-art methods such as constrained transport schemes or schemes evolving the vector potential.

We only consider evolution codes that assume the ideal MHD limit, which corresponds to the additional condition

$$E^l = -\epsilon^{ljk} v_j B_k. \quad (21)$$

This precludes the use of our scheme for resistive MHD simulations.

For convenience, we define rescaled variables as

$$\bar{q} = \frac{\bar{\tau}}{\bar{D}}, \quad \bar{r}_i = \frac{\bar{S}_i}{\bar{D}}, \quad (22)$$

$$q = \frac{\tau}{D}, \quad r_i = \frac{S_i}{D}, \quad (23)$$

$$b_i = \frac{B_i}{\sqrt{D}}, \quad b = \sqrt{b^i b_i}. \quad (24)$$

It is worth noting that for most aspects of our scheme, the relevant scale for the magnetic field is given by b , not by the commonly used magnetization, defined as the ratio between magnetic and fluid pressure. Since the pressure is typically orders of magnitude below the mass energy density, $\mathcal{O}(b) = 1$ corresponds to a very large magnetization.

We also need to decompose the momentum into parts parallel and normal to the magnetic field:

$$r_{\parallel}^i = \frac{b^l r_l}{b^2} b^i, \quad r_{\perp}^i = r^i - r_{\parallel}^i. \quad (25)$$

The decomposition is undefined for the case of zero magnetic field, but we exclusively use the product with b^2 in our scheme, which is always well defined.

D. Useful relations

In the following, we collect definitions and analytic relations for later use. First, we define two quantities that will play a central role:

$$\mu \equiv \frac{1}{Wh}, \quad x \equiv \frac{1}{1 + \mu b^2}. \quad (26)$$

Trivially, their ranges are limited to

$$0 < \mu \leq 1/h_0, \quad 0 < x \leq 1. \quad (27)$$

Given the conserved variables and μ , one can directly express the primitive variables analytically. Since the system is overdetermined, there are different possible expressions which may disagree if the given μ is inconsistent with the conserved variables. In the latter case, some expressions can diverge. We will use the ambiguity to cast the recovery into a root-finding problem by expressing the same variables in different ways that only agree for the correct value of μ .

As an intermediate step, we first remove the electromagnetic part of the conserved variables. From Eq. (21),

$$E^2 = x^2 \mu^2 B^2 r_{\perp}^2, \quad (28)$$

$$\bar{r}^i = x r_{\perp}^i + r_{\parallel}^i. \quad (29)$$

This allows us to compute the pure fluid part of the conserved variables. The corresponding quantities relevant for our purpose can be written as

$$\bar{r}^2 = x^2 r_{\perp}^2 + r_{\parallel}^2, \quad (30)$$

$$\bar{q} = q - \frac{1}{2} b^2 - \frac{1}{2} \mu^2 x^2 b^2 r_{\perp}^2. \quad (31)$$

We can now express the velocity as $v = \mu \bar{r}$. This expression does not guarantee, however, that $v < 1$ for any μ . One way to avoid exceeding the speed of light is to use the quantity z , which yields

$$v(z) = \frac{z}{\sqrt{1+z^2}} < 1. \quad (32)$$

Although we do not have a closed-form expression for z as a function of μ , we can use z to obtain a useful upper limit for the velocity, given by

$$z = \frac{\bar{r}}{h} \leq \frac{r}{h} \leq \frac{r}{h_0} \equiv z_0, \quad v \leq v_0 \equiv v(z_0). \quad (33)$$

After obtaining the velocity and Lorentz factor, we can extract the rest mass density and specific internal energy using the expressions

$$\rho = \bar{D}/W, \quad (34)$$

$$\epsilon = W(\bar{q} - \mu \bar{r}^2) + W - 1. \quad (35)$$

If ρ and ϵ are in the validity range of the EOS, we can now compute the pressure $P = P(\rho, \epsilon)$ and the enthalpy $h(\rho, \epsilon)$. Finally, the following expression for μ itself turns out to be useful:

$$\mu = \frac{1}{hW} = \frac{1}{hW(W^{-2} + v^2)} \quad (36)$$

$$= \frac{1}{\frac{h}{W} + \frac{v^2}{\mu}} = \frac{1}{\frac{h}{W} + \bar{r}^2 \mu}. \quad (37)$$

E. Designing the master function

In the following, we formulate the primitive variable recovery as a root-finding problem for a suitable master function. To this end, we employ the following design goals:

- (1) The function should be one-dimensional.
- (2) It should be continuous.
- (3) It should always have exactly one root, even for unphysical values of the conserved variables.
- (4) It should be well behaved in the Newtonian limit.
- (5) It should be well behaved for zero magnetic field.
- (6) There should be a known interval which contains the root, and on which the function is defined.
- (7) The root-finding procedure should not require derivatives of the EOS.

For our scheme, we use μ defined in Eq. (26) as the independent variable to solve for; i.e., we will construct a function $f(\mu)$ which crosses zero where μ is consistent with the conservative variables. The latter take on the role of fixed parameters. To construct f , we start with Eqs. (30) and (31) and define

$$\bar{r}^2(\mu) = r^2 x^2(\mu) + \mu x(\mu)(1 + x(\mu))(r^l b_l)^2, \quad (38)$$

$$\bar{q}(\mu) = q - \frac{1}{2} b^2 - \frac{1}{2} \mu^2 x^2(\mu) b^2 r_{\perp}^2. \quad (39)$$

Next, we define functions for velocity and Lorentz factor:

$$\hat{v}(\mu) = \min(\mu \bar{r}(\mu), v_0), \quad \hat{W}(\mu) = \frac{1}{\sqrt{1 - \hat{v}^2(\mu)}}, \quad (40)$$

where v_0 is the upper velocity limit from Eq. (33). Further, we define rest mass density and specific energy according to Eqs. (34) and (35):

$$\hat{\rho}(\mu) = \frac{\bar{D}}{\hat{W}(\mu)}, \quad (41)$$

$$\hat{\epsilon}(\mu) = \hat{W}(\mu)(\bar{q}(\mu) - \mu \bar{r}^2(\mu)) + \hat{v}^2(\mu) \frac{\hat{W}(\mu)^2}{1 + \hat{W}(\mu)}, \quad (42)$$

provided that the results fall within the validity range of the EOS. Otherwise, the density $\hat{\rho}$ is adjusted to the closest value within the validity range for ρ , and $\hat{\epsilon}$ to the closest value within the validity range for ϵ at adjusted density $\hat{\rho}$. In Eq. (42), we express the term $W - 1$ in a way that prevents large rounding errors in the case of small velocities. Using the EOS, we compute the pressure, defining

$$\begin{aligned}\hat{P}(\mu) &= P(\hat{\rho}(\mu), \hat{\epsilon}(\mu)), \\ \hat{a}(\mu) &= \frac{\hat{P}(\mu)}{\hat{\rho}(\mu)(1 + \hat{\epsilon}(\mu))}, \\ \hat{h}(\mu) &= h(\hat{\rho}(\mu), \hat{\epsilon}(\mu)).\end{aligned}\quad (43)$$

To close the circle, we can now express μ itself as a function $\hat{\mu}(\mu)$ based on $\hat{W}(\mu)\hat{h}(\mu)$. However, we find that this straightforward choice does not yield a function respecting our design goals. One reason is that under extreme conditions, the strong limitations introduced for \hat{P} , $\hat{\epsilon}$, $\hat{\rho}$, and \hat{W} can cause severe kinks in the function.

By trial and error, we find that using Eq. (37) results in a very different master function $f(\mu)$ that is well suited to our purposes. It is given by

$$f(\mu) = \mu - \hat{\mu}(\mu), \quad (44)$$

$$\hat{\mu}(\mu) = \frac{1}{\hat{\nu}(\mu) + \bar{r}^2(\mu)\mu}. \quad (45)$$

As an additional minor modification, we compute the variable $\nu \equiv h/W$ in two slightly different ways based on Eqs. (2) and (35), defining

$$\nu_A(\mu) = \frac{\hat{h}(\mu)}{\hat{W}(\mu)} = (1 + \hat{a}(\mu)) \frac{1 + \hat{\epsilon}(\mu)}{\hat{W}(\mu)}, \quad (46)$$

$$\nu_B(\mu) = (1 + \hat{a}(\mu))(1 + \bar{q}(\mu) - \mu\bar{r}^2(\mu)), \quad (47)$$

$$\hat{\nu}(\mu) = \max(\nu_A(\mu), \nu_B(\mu)). \quad (48)$$

The motivation behind the second form ν_B is to reduce the kink introduced by limiting \hat{h} , only limiting \hat{a} instead, while allowing ϵ to change further. For this, $(1 + \epsilon)/\hat{W}$ is computed directly from Eq. (35), which does not involve the EOS. The reason we use the smoother ν_B only when ϵ crosses the upper limit, but not when crossing the lower limit, is that decreasing $\hat{\nu}$ increases $\hat{\mu}$ and decreases the overall slope of the master function in some regimes, which is disadvantageous with respect to ensuring uniqueness.

Examples for the master function resulting for different regimes are shown in Fig. 1. As one can see, the master function is almost linear unless the Lorentz factor or magnetic scale b are large, but it remains well behaved even then. Note that we only show the root-bracketing




FIG. 1. Master function f for different regimes, combining velocities $v = 0$ and $v = 0.99$ (labeled “vlarge”), magnetic field scale $b = 0$ and $b = 2$ (“blarge”), internal specific energy $\epsilon_{\text{th}} = 0$ (“cold”) and 10 (“hot”), where ϵ_{th} denotes the difference from the zero-temperature case. Velocities are oriented orthogonally to the magnetic field, which is the most difficult case. In the top panel, the independent variable μ is scaled to the initial root bracket μ_+ , and the function is scaled to the maximum value over the interval shown. The lower panel shows the behavior near the root $\tilde{\mu}_0$.

interval that is constructed in the next section. Beyond this interval, the function can have strong kinks.

F. Existence of solution

In the following, we prove that the master function always has a root, not just for valid evolved variables, but also for invalid ones. To this end, we construct an interval over which the master function changes sign. We start by defining an auxiliary function

$$f_a(\mu) = \mu\sqrt{h_0^2 + \bar{r}^2(\mu)} - 1. \quad (49)$$

This is a smooth function which does not require evaluation of the EOS, using only the EOS-specific global lower enthalpy bound h_0 . It is easy to show that f_a is strictly increasing and has exactly one root μ_+ in the interval $(0, h_0^{-1}]$. We will show that μ_+ constitutes an upper bound for the root of the master function f . For $f_a(\mu_+) = 0$, we find

$$\begin{aligned}\mu_+\bar{r}(\mu_+) &= \frac{\bar{r}(\mu_+)}{\sqrt{h_0^2 + \bar{r}^2(\mu_+)}} \\ &\leq \frac{r}{\sqrt{h_0^2 + r^2}} = v_0,\end{aligned}\quad (50)$$

where we use $r \geq \bar{r}$ and the monotonicity of the above expression with respect to \bar{r} . From Eq. (40), we find $\hat{v}(\mu_+) = \mu_+ \bar{r}(\mu_+) \leq v_0 < 1$. Further, $f_a(\mu_+) = 0$ implies that

$$\mu_+ h_0 = \sqrt{1 - \hat{v}^2(\mu_+)} = \hat{W}^{-1}(\mu_+). \quad (51)$$

Using the definition Eq. (48) of \hat{v} , which implies $\hat{v} \geq \nu_A$, we can write

$$\mu_+ \hat{v}(\mu_+) \geq \mu_+ \frac{\hat{h}(\mu_+)}{\hat{W}(\mu_+)} \geq \mu_+ \frac{h_0}{\hat{W}(\mu_+)}, \quad (52)$$

$$= \hat{W}^{-2}(\mu_+) = 1 - \mu_+^2 \bar{r}^2(\mu_+). \quad (53)$$

Hence,

$$1 \leq \mu_+ (\hat{v}(\mu_+) + \bar{r}^2(\mu_+) \mu_+) = \mu_+ / \hat{\mu}(\mu_+). \quad (54)$$

Using the definition of the master function, Eq. (44), we find that $f(\mu_+) \geq 0$, and, trivially, $f(0) < 0$. Since f is continuous, it has at least one root in the interval $(0, \mu_+]$. From Eq. (52), we also find that the root is strictly below μ_+ unless $\hat{h}(\mu_+) = h_0$.

Conveniently, the interval provides a useful initial bracketing for the root-finding algorithm. Although finding μ_+ requires another numerical root solving, the computation of f_a does not require the expensive evaluation of the EOS. Moreover, determining μ_+ is not required if $r < h_0$. In this case, $\hat{v}(1/h_0) < 1$ and $f_a(1/h_0) > 0$, hence one can safely use $(0, 1/h_0]$ to bracket the root.

The main reason to expend this effort is to ensure that $\hat{v} \leq v_0 < 1$. Beyond μ_+ , the cutoff in Eq. (40) can induce a strong kink in the master function, reducing efficiency of the main root finding. With the tight upper limit μ_+ , the only reason for the cutoff is to make absolutely certain that not even rounding errors in ultra-relativistic cases can lead to not-a-number results. Finally, being able to use $\hat{v} \leq v_0$ simplifies the uniqueness proof in the next section.

G. Uniqueness of solution

Uniqueness of physically valid solutions is obviously important for any evolution scheme based on the conserved variables considered in this work. For the purpose of our recovery scheme, we require in addition that (i) for valid evolved variables, the master function has no additional roots corresponding to invalid solutions, and (ii) for invalid evolved variables, it still has exactly one root. In this section, we will prove that the above conditions are met.

We first compute the derivative of the master function. Differentiation and straightforward algebraic computations yield

$$\frac{dx}{d\mu} = -x^2 b^2, \quad (55)$$

$$\frac{d\bar{r}}{d\mu} = -\frac{(1-x)x^2}{\mu\bar{r}} r_{\perp}^2, \quad (56)$$

$$\frac{d\bar{q}}{d\mu} = -(1-x)x^2 r_{\perp}^2, \quad (57)$$

$$\frac{d}{d\mu}(\mu\bar{r}(\mu)) = \frac{1}{\bar{r}}(x^3 r_{\perp}^2 + r_{\parallel}^2) \geq 0. \quad (58)$$

Since $\mu\bar{r}(\mu)$ is monotonically increasing and we have shown in Sec. II F that $\mu_+ \bar{r}(\mu_+) \leq v_0$, it follows that for $\mu \leq \mu_+$, Eq. (40) reduces to $\hat{v}(\mu) = \mu\bar{r}(\mu)$. From this, we find

$$\frac{d}{d\mu} \ln(\hat{W}) = \hat{W}^2 \mu (x^3 r_{\perp}^2 + r_{\parallel}^2). \quad (59)$$

In the following, we assume that the density \bar{D}/\hat{W} does not leave the allowed range of the EOS. This corner case is discussed in the Appendix. We then obtain

$$\frac{d}{d\mu} \ln(\hat{\rho}) = -\frac{d}{d\mu} \ln(\hat{W}). \quad (60)$$

So far, we have computed derivatives of quantities that do not depend on Eq. (42), and therefore we did not need to consider the limiting of ϵ to the valid EOS range. For the derivative of $\hat{\epsilon}$, we first consider the case in which ϵ computed by Eq. (42) is in the valid range. We then find

$$\frac{d}{d\mu} \hat{\epsilon} = \left(1 + \hat{\epsilon} - \frac{1}{\hat{W}\mu}\right) \frac{d}{d\mu} \ln(\hat{W}). \quad (61)$$

At a solution, $\mu\hat{W}\hat{h} = 1$, and we obtain

$$\frac{d}{d\mu} \hat{\epsilon} = -\frac{\hat{P}}{\hat{\rho}} \frac{d}{d\mu} \ln(\hat{W}) = \frac{\hat{P}}{\hat{\rho}^2} \frac{d}{d\mu} \hat{\rho}. \quad (62)$$

This means that changes of density $\hat{\rho}$ and specific energy $\hat{\epsilon}$ are adiabatic when varying μ near a solution; i.e., the derivative of specific entropy \hat{s} is zero. For the case in which ϵ computed by Eq. (42) is below the valid range of the EOS, $\hat{\epsilon}$ is set to the lower bound of the validity range, $\epsilon_{\min}(\hat{\rho})$, which is the zero-temperature limit according to our EOS requirements. Consequently, $\hat{\rho}(\mu)$ and $\hat{\epsilon}(\mu)$ follow a curve of constant s . In both of the above cases we can therefore compute the derivative of \hat{h} using the adiabatic sound speed given by Eq. (8) and the derivative of density given by Eq. (60), obtaining

$$\frac{d}{d\mu} \ln \hat{h} = -\hat{c}_s^2 \frac{d}{d\mu} \ln \hat{W}, \quad (63)$$

where $\hat{c}_s = c_s(\hat{\rho}, \hat{\epsilon})$. In both cases, $\hat{v} = \hat{v}_A = \hat{h}/\hat{W}$, and we find

$$\frac{d}{d\mu} \ln \hat{v} = -(1 + \hat{c}_s^2) \frac{d}{d\mu} \ln \hat{W}. \quad (64)$$

At a solution, the derivative of the master function becomes

$$\begin{aligned} \frac{d}{d\mu} f(\mu) &= 1 - \hat{v}^2 + \hat{v}^2(1 - c_s^2) \frac{x^3 r_{\perp}^2 + r_{\parallel}^2}{\bar{r}^2} \\ &\geq 1 - \hat{v}^2 > 0, \end{aligned} \quad (65)$$

as shown in the Appendix. The requirement $c_s < 1$ is therefore sufficient to guarantee uniqueness of the root for all velocities and magnetic fields. Note that for superluminal sound speeds, we cannot prove uniqueness unless $v^2 c_s^2 < 1$. In that case, the proof of existence still holds, but is not known if the physical solution becomes ambiguous.

So far, we have not addressed the corner case where the specific energy is above the valid range of the EOS. In that case, \hat{v} is computed from Eq. (47). A straightforward computation (see the Appendix) reveals that uniqueness is always ensured under the condition that

$$\frac{A(\rho)}{1 + \hat{a}} \frac{\partial a}{\partial \epsilon} \leq 1 - c_s^2. \quad (66)$$

The function A is defined by the relation

$$\frac{A(\rho)}{\rho} = \frac{d}{d\rho} \epsilon_{\max}(\rho) - \frac{P}{\rho^2}. \quad (67)$$

It is related to the change of specific entropy along the upper validity range ϵ_{\max} of the EOS. For $A = 0$, the equation above reduces to the thermodynamic condition for adiabatic change. The condition given by Eq. (66) does not seem to be very restrictive. If in doubt, one can always use a boundary with $A = 0$ (constant specific entropy) to guarantee uniqueness in all cases. In practice, we encountered no problems using upper validity bounds defined by either constant temperature or constant ϵ .

H. Guaranteed accuracy

Since the root of the master function is determined numerically, we require a criterion to stop the iteration once sufficient accuracy is reached. What is sufficient depends on the other errors present in a numerical evolution scheme. We will discuss evolution errors in Sec. V. In this section, we discuss the error propagation of the root-finding accuracy to quantify the accuracy of the recovered primitives.

However, we first need to specify how the final result is computed from the outcome of the last root-finding iteration. This involves a design decision, since the

available variables $\mu, \hat{\mu}, \hat{v}, \bar{r}, \bar{q}, \hat{v}, \hat{W}, \hat{\rho}, \hat{\epsilon}, \hat{h}, \hat{P}$ allow us to compute the primitives in many different ways, which lead to different error propagation. Here, we use $\hat{W}, \hat{\epsilon}, \hat{\rho}, \hat{P}$ directly, which turns out to be a good choice in terms of error propagation. To reconstruct the velocity vector, we use the expression

$$\hat{v}^i = \mu \bar{r}^i = \mu x(r^i + \mu(b^l r_l) b^i), \quad (68)$$

which is just Eq. (29) rearranged to avoid degeneracy for the case $b = 0$. It is easy to verify that the Lorentz factor $W(\hat{v}^2)$ corresponding to Eq. (68) is exactly $W(\hat{v}^2) = \hat{W}$. Since $\hat{\rho} = D/\hat{W}$, the conserved density $\hat{D} = \hat{\rho} W(\hat{v}^2)$ computed from the recovered primitive variables $\hat{\rho}, \hat{v}^i$ agrees exactly with the original one.

In the following, we only consider the case where the solution is in the validity region of the EOS. For invalid solutions, the accuracy of the solution is less relevant, since in this case the cause is the evolution error, and the result will either be corrected to the valid range or the simulation aborted. The error introduced by such corrections will be discussed in Sec. V D.

Assuming the root of $f(\mu)$ was determined numerically to an accuracy of $\delta\mu$, we now estimate the resulting accuracy of the primitive variables to linear order, computing—e.g., $\delta\hat{W} = \delta\mu d\hat{W}/d\mu$. We already evaluated the first derivatives at a root of f in Sec. II G. From those, we obtain

$$\frac{\delta\hat{W}}{\hat{W}} \leq \hat{v}^2 \Delta, \quad \Delta \equiv \hat{W}^2 \frac{\delta\mu}{\mu}, \quad (69)$$

$$\frac{\delta\hat{z}}{\hat{z}} \leq \Delta, \quad \frac{\delta\hat{v}}{\hat{v}} \leq \frac{|\delta\hat{v}^i|}{\hat{v}} \leq \frac{\Delta}{\hat{W}^2}, \quad (70)$$

$$\frac{\delta\hat{\rho}}{\hat{\rho}} \leq \hat{v}^2 \Delta, \quad \frac{\delta\hat{h}}{\hat{h}} \leq \hat{v}^2 \Delta, \quad (71)$$

$$\frac{\delta\hat{\epsilon}}{1 + \hat{\epsilon}} \leq \hat{a} \hat{v}^2 \Delta, \quad \frac{\delta\hat{\epsilon}}{\hat{\epsilon}} \leq (1 + \hat{\epsilon}) \frac{\hat{a}}{\hat{\epsilon}} \hat{v}^2 \Delta, \quad (72)$$

$$\frac{\delta\hat{\rho}_E}{\hat{\rho}_E} \leq 2\Delta, \quad \frac{\delta\hat{P}}{\hat{P}} \leq \hat{v}^2 (1 + \hat{a}) \frac{\hat{c}_s^2}{\hat{a}} \Delta, \quad (73)$$

where $|\delta\hat{v}^i|$ denotes the norm given by the 3-metric of the vector δv^i . The error in the recovered primitive variables corresponds to errors $\delta q, \delta S_i$ of conserved variables. We estimate these errors to linear order, by inserting $\hat{\rho}, \hat{h}, \hat{P}, \hat{v}^i$ into Eqs. (15)–(20) and (21), and then evaluate the first derivatives with respect to μ at the root. We obtain the following scaling:

$$\frac{|\delta S_i|}{|S_i|} \leq \Delta, \quad \frac{\delta(1+q)}{1+q} \leq 4v^2 \Delta. \quad (74)$$

We find that the accuracy in μ required for a fixed relative error of the primitives increases with increasingly relativistic velocities. On the other hand, the magnetic scale b has no impact on the error bounds. It is also worth noting that the error δs of the specific entropy s is zero (to linear order in $\delta\mu$) because the variation of $\hat{\rho}$, $\hat{\epsilon}$ with respect to μ is adiabatic (see Sec. II G). Finally, we note that the above error bounds do not include numerical rounding errors. Those will be discussed in Sec. IV B.

III. ENFORCING VALIDITY

In typical numerical simulations, the evolved magnetohydrodynamic variables frequently reach an invalid state at some points, mainly due to ordinary numerical error, but also external influences such as gauge pathologies near the centers of newly formed black holes. Often, such violations are harmless and can be corrected. Any such correction turns unphysical conditions into regular evolution errors, and obviously different prescriptions will lead to different errors, both in magnitude and in character. Although correcting violations should be regarded as part of the evolution scheme, some basic pointwise corrections can be incorporated into the primitive recovery code, granting it power to change the evolved variables. The following effects cause typical harmless violations:

- (1) When evolving zero-temperature initial data, arbitrary small evolution errors can lead to evolved variables that correspond to a fluid energy density below the zero-temperature limit.
- (2) At numerical grid points at the surface of neutron stars moving through vacuum, mass and energy densities during a single time step can drop by orders of magnitude or even become negative. Although the absolute errors of the conserved variables remain small compared to the global scales of the system, the resulting local error of the specific internal energy and velocity can become huge and lead to an invalid state. The effect is alleviated over time because the errors tend to heat the outermost layer of NS surface, creating a hot atmosphere that reduces the density gradient.
- (3) During collapse to a black hole, mass density and/or temperature might leave the range covered by the given EOS, arriving at a state that is not unphysical but cannot be evolved further. This typically occurs in regions already inside the horizon or about to be engulfed by a rapidly expanding apparent horizon.
- (4) The coordinates near a black hole center are strongly stretched for gauges like the puncture gauge, and the surroundings are extremely under-resolved numerically. Under those conditions, all kinds of numerical instabilities can occur for the combined magnetohydrodynamical and spacetime evolution system.

A. Simple corrections

By design, our primitive variable recovery scheme is able to deal also with invalid input. As a side effect, we obtain a projection onto the valid regime by simply recomputing the evolved variables from the recovered primitives. As described in Sec. II E, the scheme always yields a pair $\hat{\rho}$, $\hat{\epsilon}$ such that $\hat{\epsilon}$ is within the validity range of the EOS at $\hat{\rho}$.

We first consider the important case in which the raw value of ϵ is below the valid range. In this case, only the recomputed conserved energy τ changes, while S and D stay the same. This can be seen as follows. The only variable through which the adjustment of ϵ to the valid range impacts the master function is $\hat{\nu}$. For the case at hand, Eq. (48) implies $\hat{\nu} = \nu_A$. Furthermore, the conserved energy τ enters exclusively through Eq. (42). Therefore, if τ is adjusted such that Eq. (42) yields the range-limited value for $\hat{\epsilon}$, we arrive at the same primitive variables without adjustment.

For the case in which the energy is above the validity range of the EOS, all recomputed conserved variables can change. One could prevent this by always using $\hat{\nu} = \nu_A$, but not without changing the behavior of the master function away from the solution. However, this case is less important, because this correction should only be allowed at low-density fluid-vacuum boundaries (NS surfaces) or inside horizons.

In the interiors of black holes, it becomes necessary to employ a more lenient error policy than outside. Although physical effects cannot propagate out of the horizon, violations of the constraint equations and gauge effects impact the exterior. Therefore, one cannot allow any runaway instability inside the horizon. For the matter part, this mainly concerns energy and momentum, since the total baryon number is conserved in finite-volume schemes (artificial atmosphere corrections aside), and the mass density remains finite. The energy can be limited by allowing the aforementioned correction to the EOS range inside horizons even at high densities.

This leaves the momentum. For pure hydrodynamic simulations, limiting the velocity proved effective to prevent runaway instabilities near the BH center. This was employed for the simulations in Ref. [23], by rescaling the velocity to stay within a given limit. For MHD simulations, this approach has a side effect. Since the reconstructed electric field depends on the velocity via Eq. (21), it will also change. That might be problematic or not, depending on the evolution scheme. The evolution of the EM field might be problematic in this regime in any case. However, addressing such problems is clearly not inside the scope of the primitive variable recovery, since it operates pointwise and cannot change electric or magnetic fields in any reasonable way.

Another correction often applied is to enforce a minimum mass density, also called artificial atmosphere. There are two motives. One is the wish to use a tabulated EOS

that does not include zero density (this might be achieved more consistently by extending the range to zero via analytic expressions). The more fundamental motive is that the hydrodynamic evolution equations break down in vacuum. In purely hydrodynamic simulations, it is common to set the atmosphere velocity to zero with respect to the simulation's coordinate system, in order to prevent an unphysical influx of matter. In ideal MHD simulations, the situation is more complicated because the electric field is tied to the velocity via Eq. (21). Therefore, the atmosphere velocity prescription should be the domain of the evolution scheme and not of the primitive recovery.

IV. PERFORMANCE

In the following, we assess how well our scheme performs in practice. For this, we subject a reference implementation to a series of numerical tests. These stand-alone tests simulate conditions relevant for the future use of the scheme in actual numerical simulation codes.

Our tests aim to validate that the following requirements are met. First, the scheme should not fail to converge for any input we expect to occur in numerical simulations of binary neutron star mergers (or supernova core collapse, which has similar demands). Second, within the physical regimes that may occur in the above scenarios, the errors caused by the primitive recovery should be insignificant compared to the ones caused by modern evolution schemes.

Lorentz factor W , magnetic scale b , and specific energy ϵ are the most important variables governing the behavior of the primitive recovery. To get an indication of what values to expect for b , we consider the merger simulations described in Refs. [37,38] as a typical use case. The maximum magnetic field strength reached after the merger by various amplification mechanisms is $\sim 10^{17}$ G. The lowest density reached in the funnel along the axis after black hole formation is around $\sim 10^8$ g/cm³, although these conditions do not coincide.

The corresponding value $b \approx 100$ constitutes a reasonable scale up to which we demand robustness. A comparable value $b \approx 30$ corresponds to initial data with neutron stars possessing a magnetar-like exterior field strength 10^{15} G that extends into an artificial atmosphere with typical densities of 10^5 g/cm³. In merger simulations, the regions with magnetic fields amplified even further do not overlap with the lowest-density regions. Combining a field $\sim 10^{17}$ G with the aforementioned artificial atmosphere density, we find that $b \approx 10^3$ exceeds practical use cases by far, and demanding robustness up to this scale is not required.

Although the typical Lorentz factors we encounter in merger simulations are well below 10, one motivation for magnetized merger simulation is to observe the launch of a jet, and future applications might follow up such a jet to higher Lorentz factors. High Lorentz factors might also

appear inside black holes. Therefore, we demand robustness up to $W = 10^3$.

Regarding the requirements for specific energy, we note that nuclear matter at the highest densities possible in NSs can reach $\mathcal{O}(\epsilon) = 1$ at zero temperature. Although temperatures can easily reach ~ 50 MeV inside merging NSs (see Ref. [39], for example), this happens at high densities. The thermal contribution to ϵ stays well below unity (see, e.g., the appendix of Ref. [40]). Furthermore, for states with $\epsilon \gg 1$, the energy of the photon gas dominates that of baryonic matter, and the evolution equations of magneto-hydrodynamics are not applicable anymore. It seems reasonable to demand accuracy and robustness of the primitive recovery up to $\epsilon \approx 10$.

In order to understand finite precision effects, we compare the accuracy that our reference implementation reaches to the expected accuracy derived in Sec. II H, using a value $\Delta = 10^{-8}$. We assume that such an accuracy is negligible compared to the evolution error. Note, however, that in Sec. V, we show that the accuracy of ϵ (and hence P) related to the accuracy of the evolved variables deteriorates quickly for $b > 1$, with a factor $\approx b^2/\epsilon$. We therefore restrict tests of primitive recovery accuracy to the regime $b < 5$ and caution that one cannot trust the evolution scheme at higher values.

A. Code design

We created a reference implementation for the algorithm described in Sec. II. A schematic summary of the algorithm is given in Fig. 2. We provide our reference implementation in the form of a C++ library named `RePrimAnd`. The library is not tied to any particular evolution framework, allowing the use in arbitrary evolution codes. It also contains a framework providing access to different types of EOS through a generic interface, ensuring that the user code (such as our primitive recovery code) is completely EOS-agnostic. The generic interface also provides the EOS validity ranges and rigidly enforces our EOS requirements (see Sec. II B). The reference implementation is publicly available [34].

Our library provides different EOS types (not yet including fully tabulated ones), including a hybrid EOS based on a tabulated cold part. Although the MS1 variant from Ref. [35] used for our test is given analytically in the form of piecewise polytropic expressions, we evaluate it in our tests as a tabulated cold part, in order to test the general-purpose code intended for production runs. We set the allowed range of the EOS to $\rho_{\max} = 3 \times 10^{15}$ g/cm³, $\epsilon_{\max} = 51$.

In order to find the root of the master function, our implementation uses the TOM748 algorithm [41] provided by the BOOST library. This root-finding scheme is similar to the well known Brent-Dekker schemes, but it uses inverse cubic instead of quadratic interpolation whenever possible, improving convergence speed near the solution. It keeps the solution enclosed in a bracket, with extent




FIG. 2. Schematic overview of the recovery algorithm information flow and list of required equations. Arrows denote dependencies, including variables further up in the chain. Double arrows denote computations that require information about the EOS validity bounds; the one double arrow labeled “EOS” refers to evaluation of the EOS. Intermediate results obtained during evaluation of the master function in the last root finding iteration are denoted “Last call.” Square brackets refer to the list of fixed parameters (independent of μ) during root solving. The density ρ_{bnd} stands for any one of the upper and lower EOS validity bounds.

converging in a limited amount of steps, and it does not make use of function derivatives. The motivation for avoiding derivatives is that, in practice, tabulated EOSs tend to have very inaccurate partial derivatives, which is problematic when using a derivative-based root solver such as Newton-Raphson. Our implementation, therefore, does not make use of the sound speed or other derivatives. The root is determined to an accuracy specified by prescribing Δ defined in Eq. (69), where the error $\delta\mu$ is taken as the size of the current tightest bracket of the root.

In the case where $r \geq h_0$, we need to determine μ_+ from the root of the auxiliary function f_a , for which we employ a standard Newton-Raphson root solver. This is unproblematic, since f_a is a smooth, monotonic, analytic function. We determine μ_+ to an accuracy close to machine precision, and then increase μ_+ by a multiple of the root-solving accuracy to ensure the root of the master function is really contained.

B. Robustness and accuracy

Our main test validates both robustness and expected accuracy by sampling the primitive variable parameter

space given by density, temperature/specific energy, magnetic scale b , and velocity. We sample $z = Wv$ between 0 and 1000, magnetic scale b from 0 to 5, and the specific thermal energy from $\epsilon_{\text{th}} = 10^{-4}$ up to 50. For the MS1 EOS, we sample the mass density from 10^6 to 10^{15} g/cm³. For the ideal gas EOS, the mass density is irrelevant due to the scaling behavior of the EOS. We use two orientations of the velocity, parallel and orthogonal to the magnetic field. The tests are performed both for the ideal gas and for the hybrid EOS, described in Sec. II B, and we demand an accuracy $\Delta = 10^{-8}$.

We verify that the algorithm always, without any exception, succeeds in recovering the correct solution for the valid input described above. Moreover, we create test cases to assure that input corresponding to energy outside the range possible for a given EOS is correctly classified as such.

To assess the accuracy, we compute the conserved variables from the primitives, apply the primitive recovery algorithm, and compare the result to the original primitives. Further, we compute the conservatives from the recovered primitives and compare them to the original conservatives. Our test suite compares the observed accuracy for each individual primitive variable to the one expected from Eqs. (69)–(73), and also the accuracy of the corresponding conserved variables to Eq. (74). When demanding an accuracy better than $\Delta \lesssim 10^{-7}$, those bounds are exceeded either for high Lorentz factors or very small ϵ and v . We attribute the excess error to various rounding errors.

We identify the most important rounding errors as follows. First, the master function is the difference of two values which can each be expressed only to machine precision. To get the impact on the root, we have to divide by the derivative of the master function, which in this case satisfies $f' \leq 1 - v^2 c_s^2$. At the same time, we demand an accuracy Δ/W^2 . For the highly relativistic case $W = 10^3$ and around 16 digits machine precision, this limits $\Delta > 10^{-10}$. If the sound speed approaches unity, the accuracy is further limited. Second, ϵ is computed by subtracting kinetic and magnetic energy density from the total one. If ϵ is small compared to these, the cancellation error causes a loss of significant digits. Analyzing Eq. (42), we find additional rounding errors of magnitudes z^2/ϵ and $b^2 W/\epsilon$ worse than machine precision.

Taking into account both the regular errors predicted by Eqs. (69)–(73) as well as the main rounding errors discussed above, the recovered accuracy is quantitatively within the expected bounds over the whole range of our test cases. Figure 3 shows the recovered accuracy for the pressure as well as the boundary where the errors caused by rounding start exceeding those caused by root finding.

We do not expect rounding errors to be of practical importance. The rounding errors at low ϵ , v are very small and only dominate because the regular errors approach zero. The rounding errors in the ultrarelativistic/highly




FIG. 3. Relative error of reconstructed pressure, as a function of specific thermal energy and velocity (the latter in terms of $z = Wv$). The results were obtained for the case of the hybrid EOS (see text), at fixed mass density 6×10^{12} g/cm³, demanding an accuracy $\Delta = 10^{-8}$. The upper panel shows results for zero magnetic field, and the lower panel for the magnetically dominated case $b = 10$. The solid lines mark the regions where expected errors related to rounding start to exceed those related to root solving.

magnetized regime are still not prohibitive, but will be dominated by the errors of the time evolution, which will be discussed in Sec. V.

C. Efficiency

In the following, we discuss the efficiency of our scheme on the level of the algorithm, while reserving benchmarks of the execution speed of the implementation within actual GRMHD simulations for future work. We measure the computational efficiency of our algorithm in terms of calls to the EOS. The motivation is that for a tabulated EOS including thermal and composition degrees of freedom, a single EOS call is likely more expensive than the evaluation of the analytic expressions within our recovery scheme. The worst scenario is when the EOS is tabulated with temperature as one independent variable. Each EOS call then requires an inversion step to convert from ϵ to T .

Figure 4 shows how the efficiency varies with specific energy and velocity, either for zero magnetic field, or with magnetic scale fixed to a large value of $b = 10$. We find that the efficiency does not degrade even for Lorentz factors up to 1000 and magnetic scales up to $b = 10$. At the density shown, $b = 10$ corresponds to extremely high magnetizations of orders 10^4 (for $W = 1$) to 10^7 (for $W = 1000$).

When considering the whole parameter space used in the unit tests (not just the cuts shown in the plots) and both EOS types, we find a maximum number of 23 calls to the EOS required to achieve an accuracy $\Delta = 10^{-8}$. The maximum occurs for the ideal gas and only when both $\epsilon > 40$ and $b > 2$ —i.e., thermal energies much larger and magnetic energies larger than the rest mass density.

In Fig. 5, we show the efficiency with respect to velocity and magnetic scale b , taking the latter up to extreme values




FIG. 4. Number of calls to the EOS required to reconstruct the primitives to accuracy $\Delta = 10^{-8}$, as a function of specific thermal energy and velocity (in terms of $z = Wv$). The results were obtained for the case of the hybrid EOS (see text) at a mass density $\rho = 6 \times 10^{12}$ g/cm³. The upper panel shows results for the magnetic scale $b = 0$, and the lower panel for the magnetically dominated case $b = 10$.

$b = 10^4$, far beyond any reasonable use case. We find that beyond $b = 10$, the efficiency gradually starts to decrease. At $b = 10^4$, $W = 10^3$, we require around 40 steps for $\Delta = 10^{-8}$ (which implies $\delta\mu/\mu \approx 10^{-14}$). At this point, the root-solving convergence speed has decreased roughly to that of bisection. Still, we encounter no failures to converge even in this range.

Note that the extremes reached in our tests are rather pathological scenarios which are rarely encountered in simulations and are therefore not relevant for numerical costs of simulations. In practice, we expect an average number of required calls below 10.




FIG. 5. Number of calls to the EOS required to reconstruct the primitives to accuracy $\Delta = 10^{-8}$, as a function of magnetic scale b and velocity (in terms of $z = Wv$). The results were obtained for the case of the hybrid EOS (see text) at a mass density $\rho = 6 \times 10^{12}$ g/cm³. The upper panel shows results for cold matter $\epsilon_{\text{th}} = 10^{-4}$, and the lower panel for very hot matter $\epsilon_{\text{th}} = 10$. For comparison, we also show the magnetization as contour lines of $\log_{10}(B^2/P)$. The red line marks a magnetic field strength of $B = 10^{16}$ G.

TABLE I. Main characteristics of different recovery schemes. We list the independent variables used in the root finding (translated to our notation), the variables for which the EOS needs to be evaluated, whether the scheme requires derivatives of the EOS, whether the formulation allows a bound on the number of iterations needed for finding the solution, and whether the scheme requires one to provide an initial guess for the solution.

Scheme	Independent variables	EOS form	EOS derivatives	Steps bounded	Guess needed
This work	μ	$P(\rho, \epsilon)$	No	Yes	No
Noble [30]	$(D/\mu, v^2)$	$P(\rho, h)$	Yes	No	Yes
Siegel [33]	$(D/\mu, T)$	$P(\rho, T)$	Yes	No	Yes
Duran [26]	$(W, D/\mu, T)$	$\epsilon(\rho, T)$	Yes	No	Yes
Neilsen [31,32]	D/μ	$P(\rho, \epsilon)$	No	Yes	No
Newman [42]	P	$P(\rho, h)$	No	Yes	No

D. Comparison with other schemes

In the following, we compare our scheme to existing ones. We refer to Ref. [33] for a comprehensive review and numerical tests of previous schemes. The main characteristics are listed in Table I.

One important difference is the number of independent variables. Most of the existing schemes need to solve an equation in two or three unknowns. This is a severe drawback. First, it is difficult to ensure that the solution is found. The Newton-Raphson (NR) schemes might not converge. Second, robust but fast schemes that guarantee finding the solution in a limited number of steps only exist for one-dimensional root finding. Third, the recovery schemes based on NR require an initial guess, which is typically taken from the previous time step during numerical evolution. This makes the methods more unpredictable and more difficult to test, as they do not depend on the conserved variables in a deterministic way. As two of the existing schemes, our scheme is using one-dimensional root finding. Further, it also makes use of a tight initial bracketing interval proven to contain exactly one solution.

As demonstrated in Ref. [33], all of the existing schemes can fail for Lorentz factors 10–1000, depending on the magnetization. Figure 3 of Ref. [33] shows the number of iterations or failure to converge as a function of magnetization and Lorentz factor, at fixed density 10^{11} g/cm³ and $T = 5$ MeV (thus $\epsilon < 1$). For our scheme, b is the most relevant measure for the magnetic field (but not necessarily for the other schemes). For comparison, the magnetization 10^{10} covered in the figure corresponds to values up to $b \approx 10^4$. As argued before, such values are outside the parameter space relevant for use in merger simulations. Therefore, the failures at low velocity, but with magnetization around 10^9 shown in Fig. 3 of Ref. [33] are not problematic in practice. The fact that our scheme showed

no failure at $b = 10^4$ for the test shown in Fig. 5 is, however, reassuring regarding the numerical robustness. The failures at relativistic velocities for lower magnetization shown in Fig. 3 of Ref. [33] should, however, be regarded as problematic. For our algorithm, the existence and uniqueness of the solution are proven analytically, and we successfully test our numerical implementation up to Lorentz factors $W = 1000$ in the whole parameter space described in Sec. IV B.

Strong magnetization is important for studying the engine of short gamma-ray bursts (SGRBs), where ideally a very low-density matter is subject to very strong magnetic fields. This regime is also problematic for the numerical time evolution itself. The ability of our scheme to distinguish reliably between valid and invalid evolved variables is therefore an important advantage.

Figure 2 of Ref. [33] shows the average of the relative errors (called $\bar{\sigma}$ in the following) for ρ, v^i, ϵ as a function of temperature and density for fixed low magnetization $p_{\text{mag}}/p = 10^{-3}$ and Lorentz factor $W = 2$. We note that for the nuclear physics EOS, one can ignore the top-left part of the plots in Fig. 2, because this low-density, high-temperature regime corresponds essentially to a photon gas not relevant for practical use. Using the energy density $\rho_\gamma \propto T^4$ of a photon gas, we find that the bound $\epsilon < 10$ used in our tests corresponds to temperatures below a straight line (in the log-log plot) through $(10^{12}$ g/cm³, 10^{12} K) and $(10^4$ g/cm³, 10^{10} K). For the ideal gas EOS, the temperature is instead defined assuming only baryons, such that $\epsilon < 1$ throughout the figure.

In most of the physically relevant region, the recovery accuracy shown in Ref. [33] seems more than sufficient for use in evolution schemes. However, some existing schemes exhibit isolated regions where the accuracy degrades inexplicably. This is worrisome, since the figure shows only a mildly relativistic two-dimensional cut in the parameter space, and there is no guarantee that the accuracy will not deteriorate intolerably elsewhere. In contrast, our scheme has the advantage of a theoretical model for the errors of each of the primitive variables, including the dominant finite-precision errors. This model was validated for our numerical implementation over the full parameter space (ρ, W, ϵ, b) as described in Sec. IV B.

Regarding the efficiency, the different tolerance measures allow only a rough comparison, using the number of root-finding iterations shown in Figs. 1 and 3 of Ref. [33]. Note that the top-left part of the plots in Fig. 1 is not practically relevant for the nuclear physics EOS, as discussed above. As mentioned before, the problems at the highest magnetizations $> 10^9$ in Fig. 3 of Ref. [33] can be safely ignored for practical use. In case no failure occurs, only the Newman scheme appears to be consistently requiring less than around 10 steps also for relativistic velocities. The others need 30 iterations or more in certain regions of parameter space.

Our scheme is guaranteed to converge in a finite number of steps, because the root-finding algorithm performs bisection steps if needed. Our tests have shown that the efficiency does not degrade for large Lorentz factors ($W < 1000$) or strong magnetization ($b < 100$), in contrast to most other schemes. The worst-case scenario for our scheme seems to be extreme values of ϵ . Even for essentially photonic states ($\epsilon = 50$), it does not require more than 23 EOS calls in the regime $b < 5$, $W < 1000$.

Up to this point, we considered the convergence criteria used during the iteration as an integral part of the different algorithms. Now we have to discuss if the different measures of success could bias the comparison. In Ref. [33], recovery was called successful when an average error $\bar{\sigma} < 5 \times 10^{-8}$ could be achieved. Comparing this to the root-solving errors given by Eqs. (69)–(73), we find that $\bar{\sigma} < K\Delta$, where K is a constant of order unity.¹ However, the term $\delta\epsilon/\epsilon$ is also very sensitive to unavoidable cancellation errors, as already discussed in Sec. IV B. Taking into account the finite floating-point precision of q , r , b , we find that catastrophic cancellation decreases the number of valid digits for ϵ by around $\log_{10}(z^2/\epsilon)$ for the case $b = 0$, $z^2 \gg \epsilon$, and by $\log_{10}(b^2/\epsilon)$ for $v = 0$, $b^2 \gg \epsilon$. For the ideal gas case shown in the upper panel of Fig. 3 in Ref. [33], ϵ and P are constant, and one can easily estimate the loss of precision. Assuming a typical machine precision of 15 digits for q , r , b , rounding errors will make it impossible to reach the tolerance $\bar{\sigma}$ required in Ref. [33] at $W \gtrsim 2 \times 10^3$ (for small B) or at $P/P_{\text{mag}} \gtrsim 5 \times 10^8$ (for small v). The above errors are a fundamental consequence of evolving the conserved variables, independent of the recovery scheme. Some, if not all of the “failures” near the upper boundary of the ideal gas panels are therefore inevitable also for our scheme.

We test if our implementation can reach the $\bar{\sigma}$ tolerance from Ref. [33] within our standard test domain introduced in Sec. IV B. In general, at $b = 0$, our theoretical order of magnitude estimate predicts a larger impact of rounding errors for $\epsilon \lesssim 0.005(W/10^3)^2$. For the ideal gas, the tolerance $\bar{\sigma}$ was indeed only reached when restricting ϵ above this estimate. Apart from this regime, the tolerance was reached in the whole standard test domain ($W \leq 10^3$, $b < 5$) using $\Delta = 10^{-8}$. The same holds for the hybrid EOS. We conclude that the impact of the cancellation error on the error measure $\bar{\sigma}$ is close to the theoretical minimum for our implementation. In contrast, all schemes shown in Ref. [33] exceed the tolerance already below $W < 10^3$ at small B , where our scheme succeeds.

¹This ceases to be true near zero crossings of ϵ . Whether such crossings can occur depends on the choice of the arbitrary constant m_B (see Sec. II A). This ambiguity is a general problem with the definition of $\bar{\sigma}$, which contains a term $\delta\epsilon/\epsilon$. It is not relevant for the following discussion, but might be a pitfall in the low-density, low-temperature regime.

The above discussion indicates that $\bar{\sigma}$ might not be a good choice as a criterion for recovery failures. We also point out that the relative error of ϵ is not a good measure for the error of the temperature, which is more closely related to the thermal part ϵ_{th} . For practical applications, it seems preferable to allow the inevitable accuracy loss for ϵ discussed above, recovering each primitive variable as accurately as theoretically possible. Whether or not the unavoidable cancellation errors in ϵ are tolerable surely depends on the application and should be regarded as part of the post-recovery error policy.

As is pointed out in Ref. [33], comparing the number of EOS calls between different schemes does not directly translate to numerical costs. The reason is that the schemes differ with respect to the required form of the EOS. Our scheme is using $P = P(\rho, \epsilon)$, while nuclear physics EOS tables are typically given in terms of T instead of ϵ . It was argued in Ref. [33] that it is advantageous if the master function directly uses the temperature as the independent variable, because otherwise each call to the EOS requires another root finding to determine T . We regard this as a shortcoming of the EOS implementation and advocate against basing the design of the primitive recovery on internals of specific EOS implementations. A more natural solution is to first create new tables in terms of ϵ (or a suitable analytic function thereof) by interpolating available nuclear physics tables. This allows one to choose the most robust recovery procedure without sacrificing speed. When using a lookup table based on temperature, however, the scheme proposed in Ref. [33] might indeed be faster than ours. We point out, however, that the large speedup discussed in Ref. [33] seems to be based on a particularly wasteful implementation of the inversion $T(\epsilon)$ that can require up to 100 steps. In Ref. [23], we used a discrete bisection in index space followed by inverse interpolation, which requires < 10 steps for realistic table sizes.

In contrast to Ref. [33], we do not test the scheme with a fully tabulated EOS and can therefore make no conclusive claims on robustness and accuracy of the implementation in this case. However, the algorithm itself is guaranteed to find a solution. We recall that the proof of existence does not rely on EOS properties except for a lower bound on h . A table in conjunction with an interpolation method represents a well-defined EOS. As long as this EOS respects the physical constraints listed in Sec. II B, the uniqueness is also guaranteed. A careless implementation of a tabulated EOS that violates those constraints might, however, cause our algorithm to find wrong, unphysical solutions. For such faulty EOSs there might even exist several physical solutions. Furthermore, our accuracy bounds are not guaranteed anymore because they rely on the constraints as well. Still, we do expect the numerical root finding to converge to a solution. Even jumps in the master function would only reduce the efficiency of the TOM748 solver, possibly down to that of bisection. We recall that this

root-finding method does not use derivatives, and would therefore not be affected by an EOS implementation returning inaccurate numerical derivatives of the pressure. Such inconsistencies might affect those schemes in Ref. [33] based on Newton-Raphson root solvers. However, the failures to converge exposed in Ref. [33] also appear for the purely analytic ideal gas EOS, and can therefore not be attributed to possible faults in EOS tables.

V. IMPACT OF NUMERICAL ERROR IN EVOLVED VARIABLES

In the following, we investigate consequences of numerical errors in the evolved variables in conjunction with the corrections of invalid states described in Sec. III. Further, we identify regions in parameter space where the primitive variables are particularly sensitive to errors of the evolved ones.

A. Newtonian limit

It is instructive to consider the relation between evolved and primitive variables in the Newtonian limit. Assuming that both kinetic and thermal specific energies are non-relativistic corresponds to $v \ll 1$, $\epsilon \ll 1$, $a \ll 1$, $h \approx 1$ (for simplicity, we chose m_B such that $h_0 = 1$ in this section). To leading order in v^2 and ϵ , we obtain

$$x \rightarrow x_N = \frac{1}{1 + b^2}, \quad (75)$$

$$\bar{r}^i \rightarrow x_N r_\perp^i + r_\parallel^i, \quad (76)$$

$$v^i \rightarrow \bar{r}^i, \quad (77)$$

$$\epsilon \rightarrow q - \frac{1}{2}(b^2(1 + v_\perp^2) + v^2). \quad (78)$$

Taking the Newtonian limit locally does not imply small b . However, if the magnetic field energy is comparable to the rest mass density, one cannot expect the velocity to stay nonrelativistic during the course of the evolution. It is a plausible assumption that the density of kinetic energy is not much smaller than the magnetic energy density. Setting $\mathcal{O}(b^2) \approx \mathcal{O}(v^2)$, we find $x_N \approx 1$. Since $\mathcal{O}(E) = \mathcal{O}(vB)$, we can also neglect the electric contribution $b^2 v_\perp^2$ to ϵ .

On a side note, it is easy to show that the master function Eq. (44) becomes a linear function in the Newtonian limit (with $b^2 \ll 1$). As $x(\mu) \approx 1$, $\bar{r}(\mu)$ and $\bar{q}(\mu)$ are independent of μ and equal to the correct values. The same holds in turn for $\hat{\rho}$, $\hat{\epsilon}$, \hat{P} . Further, $\hat{v} \approx \hat{h} \approx h_0$. Since $\bar{r}^2 \ll 1$, the master function becomes $f(\mu) \approx \mu - 1$.

We now turn to the propagation of the evolution error of the variables q , r , D . Even in the Newtonian limit, both v^2 and b^2 contributions can dominate q if ϵ is even smaller. Since v is essentially computed from r , computing ϵ from

the evolved variables suffers from cancellations that amplify the evolution errors. In detail,

$$\begin{aligned} \frac{\delta\epsilon}{\epsilon} &= \frac{\delta q}{q} \left(\mathcal{O}\left(\frac{b^2}{\epsilon}\right) + \mathcal{O}\left(\frac{v^2}{\epsilon}\right) \right) \\ &+ \frac{\delta b}{b} \mathcal{O}\left(\frac{b^2}{\epsilon}\right) + \frac{|\delta r^i|}{r} \mathcal{O}\left(\frac{v^2}{\epsilon}\right). \end{aligned} \quad (79)$$

Once $\delta\epsilon/\epsilon$ exceeds unity, reconstructing ϵ from the evolved variables might lead to larger errors than simply setting it to the zero-temperature value. Assuming some bound for the relative errors of the evolved variables, this corresponds to critical values for b^2 and v^2 .

B. Magnetically dominated regime

In the context of magnetohydrodynamic evolution, ‘‘magnetically dominated’’ refers to the magnetic pressure exceeding the fluid pressure. Increasing the field strength at fixed matter density, the movement of matter becomes constrained along the field lines at some point.

This effect is also reflected in the equations we use for primitive recovery. The relation between total and fluid momentum S_\perp components perpendicular to the magnetic field is $\bar{S}_\perp = x S_\perp$, as seen from Eq. (29). The quantity x depends only on μb^2 . In the limit $\mu b^2 \gg 1$, we find $x \ll 1$. In that case, the perpendicular part of the evolved momentum is dominated by the electromagnetic part. However, the latter is proportional to v_\perp in ideal MHD, and also points in the same direction. Therefore, the evolution error of the perpendicular part of momentum is not amplified by cancellations when recovering the fluid velocity orthogonal to the field. Also, the parallel part of the evolved momentum is not problematic, since it has no electromagnetic contribution.

As discussed for the Newtonian limit, strong magnetic fields are also detrimental to the accuracy of ϵ , since evolution errors in B are amplified by cancellation. From Eq. (39), we find that the magnetic field contribution to the energy causes strong cancellation error if $\bar{q} \ll b^2(1 - v_\perp^2)$.

C. General case

To quantify the error amplification in the general case, we compute the partial derivatives of the primitive variables with respect to the conserved ones (by means of finite differences). As expected, specific internal energy and pressure exhibit large error amplification in some regimes, while the fluid momentum is well behaved. Figure 6 shows the behavior of amplification factors of the specific energy error in relation to errors in the evolution of q , b , and r , defined as

$$A_q = \frac{\delta \log(\epsilon)}{\delta \log(q)} \Big|_{D,r,b}, \quad (80)$$




FIG. 6. Amplification factors A_q , A_r , and A_b for the relative error of ϵ , as defined in Eq. (80), for the example of cold matter obeying the MS1 EOS. Top panel: amplification as a function of density, for fixed magnetic field and velocity. Bottom panel: amplification versus velocity, for fixed density and no magnetic field.

$$A_b = \left. \frac{\delta \log(\epsilon)}{\delta \log(b)} \right|_{D,q,r}, \quad (81)$$

$$A_r = \left. \frac{\delta \log(\epsilon)}{\delta \log(r)} \right|_{D,q,b}. \quad (82)$$

The error of the pressure shows the same qualitative behavior. For a magnetar-strength field $B = 10^{15}$ G, we find that a relative evolution error of $\delta b/b = 10^{-4}$ would start to dominate the evolution of ϵ at densities of magnitude 10^8 g/cm³. The same holds for a relative error $\delta q/q = 10^{-4}$, which is to be expected, since q is dominated by the b^2 contribution. This regime could be relevant for the engine of SGRBs, as a popular model assumes a low-density funnel along the rotation axis of a black hole immersed in a strong magnetic field which is anchored in a surrounding disk. Similarly, the material surrounding a supramassive (i.e., long-lived) neutron star merger remnant could be affected.

The consequence of artificial heating could be artificial outflows and increased neutrino luminosity. To assess the potential for spurious winds, we can compare the scales of additional specific energy caused by the evolution error and the specific gravitational binding energy. At a distance 100 km from a $M = 2 M_\odot$ remnant, we find that the thermal error starts to dominate at densities 10^8 g/cm³,

again for a fiducial evolution error $\delta b/b = 10^{-4}$ and $B = 10^{15}$ G.

On the other hand, the above discussion is overly pessimistic if the increase in thermal energy by physical effects, such as shocks or neutrino absorption, greatly exceeds the one by numerical errors. In other words, the presence of a mild outflow caused by mild heating should be met with more skepticism than stronger outflows caused by prominent heating.

D. Interaction with recovery corrections

Enforcing the evolved variables to stay in the physically valid regime corresponds to a projection onto the validity boundary. Depending on the choice of projection, it is possible that the corrections cause a drift along the boundary, in the worst case with a preferred direction. This is of particular concern for the very frequent correction of limiting the specific energy above the zero-temperature value. For our scheme, only the energy density is corrected in that case. Therefore, it does not introduce a drift of the evolved momentum density.

The main effect of limiting ϵ above the zero-temperature value may be to induce a spurious heating. The reason is that cutting the evolution error distribution from below creates a positive bias until the error distribution has little support below the cut. Of course, the raw error distribution before the correction could already contain a bias. For the idealized case in which the evolution errors follow a zero-mean normal distribution around the correct result, we expect the temperature to increase until the thermal energy reaches a level comparable to the width of the error distribution for total energy.

Note that excessive artificial heating could reduce the velocity, since the momentum density incorporates a factor h . However, if h is significantly increased nonhomogeneously by the errors, the corresponding changes in thermal pressure can be expected to cause gradients and corresponding acceleration.

In the above discussion, we omitted the effect of finite root-solving accuracy. Our implementation of the algorithm recomputes all conserved variables from the primitive ones if, and only if, corrections were required. Therefore, the momentum only remains constant to the accuracy of the root solving when applying the correction to the energy. This error is formally bounded by Eq. (74). We cannot predict, however, if the distribution limited by this bound is symmetric or not. Any bias might lead to a cumulative effect over many corrections. For the worst possible scenario where each correction leads to the maximum possible momentum error always pointing along the momentum, a few thousand corrections could add up to intolerable levels.

To assess the actual behavior, we performed a numerical experiment. Instead of using a full numerical evolution, we employ a random walk model representing an evolution

error, starting at selected states. After each “evolution” step, we apply the primitive recovery and limit the conserved variables to the allowed regime. The cumulative corrections applied to energy and momentum are monitored.

In this approach, we can prescribe the error distribution. As a worst-case example, we use a normal distribution with a negative mean for the energy error. Starting at a zero-temperature state, this causes frequent corrections to the energy. Note that the expected error in the momentum does not depend on the magnitude of the corrections, but this randomized test nevertheless involves different magnitudes. For the root-finding accuracy, we use four different values: $\Delta = 10^{-7}10^{-8}$, 10^{-9} , 10^{-10} .

We find that the average momentum error is orders of magnitude smaller than the limit Eq. (74). Selecting an initial state $v = 0.99$, $b = 2$, $\rho = 6 \times 10^{12}$ g/cm³, the momentum errors of individual correction steps approach machine precision levels around $\Delta = 10^{-9}$. We believe that the reason might be that the accuracy increases drastically during the final root-finding step, such that the average root error is much smaller than the prescribed maximum. We conclude that cumulative effects of the corrections can likely be neglected. In case of evidence to the contrary, the solution would be to simply not recompute the momentum, sacrificing machine-precision consistency for error reduction.

We also apply the random walk model to states with different combinations of $b = \{0, 2\}$, $v = \{0, 0.99\}$, $\epsilon_{\text{th}} = \{0, 10\}$, perturbing the evolved variables separately with normally distributed relative errors of order 10^{-4} . This test confirms that the implementation of the zero-temperature energy correction works as intended. We monitored the behavior of v^i , W , ϵ , P , μ for the above cases and did not encounter any problematic behavior.

VI. SUMMARY AND DISCUSSION

In this work, we solved the technical problem of primitive variable recovery in relativistic ideal magnetohydrodynamic evolution codes via a new fully reliable scheme. We derived a mathematical proof that the algorithm always finds a valid solution, and that the solution is unique. Moreover, we derived expressions that allow us to prescribe the accuracy of the individual primitive variables.

The guaranteed reliability of the new algorithm is a big advantage compared to older methods, which are able to handle most of the parameter space encountered in BNS merger scenarios, but may still fail in some cases Ref. [33]. Even rare recovery failures are very problematic, since they necessitate manual intervention, and may require repeating parts of the simulation. Recovery failures are practically unpredictable and potentially chaotic (we recall the Newton-Raphson fractal related to convergence properties of a standard root-finding procedure). This is aggravated for recovery schemes that rely on an initial guess taken from the previous time step. The automated approach of using a fixed state (e.g., artificial atmosphere) in case of

recovery failure will render simulations unpredictable in practice.

The ability to identify unphysical evolved variables, as well as the nature of the invalidity, is another advantage of our method. All evolution schemes produce unphysical states occasionally, most of which are harmless. However, sometimes invalid input occurs as the first symptom of more severe evolution errors. Our method allows us to prescribe an error policy and selectively apply corrections based on the nature of the problem. Such corrections (or lack of corrections) should be considered as part of the evolution scheme, but they are mentioned in the literature only rarely.

The design of our scheme naturally suggests a particular prescription for correcting slightly unphysical input. We discussed potential cumulative effects of those corrections, and predicted that it will create artificial heating if the matter is close to zero temperature. We also showed that there should be no direct impact on the momentum. We validated this by performing a numerical experiment using random walk perturbations to emulate evolution errors.

Since the implementation of recovery algorithms is a work-intensive endeavor, we are making our reference implementation public in form of a well-documented library named `RePrimAnd` [34], which can be used by any evolution code. In order to be useful in practice, the recovery should not constrain the type of EOS. Therefore, our recovery algorithm is formulated in an EOS-agnostic manner, and the reference implementation contains a generic interface for using arbitrary EOSs.

We subjected the code to a comprehensive suite of tests, demonstrating that both the algorithm and the actual implementation are robust up to Lorentz factors and values of magnetization much larger than those relevant for BNS mergers. We also showed that the scheme is computationally efficient regarding the number of EOS evaluations (efficiency of EOS implementations aside).

While investigating the accuracy of the recovery scheme, we identified regimes where rounding errors are amplified by unavoidable cancellation errors. We quantified the dominant contributions and found that the accuracy measured in our tests is compatible with the predictions. We also found that the rounding errors are irrelevant, because the very same cancellation also leads to the amplification of evolution errors. Investigating the error propagation from evolved to primitive variables, we showed that the accuracy of the thermal energy and thermal pressure can severely degrade when evolving strongly magnetized regions of low density.

We believe that our results will be useful in particular for studying the launching mechanism of jets powering SGRBs, as well as the mass ejection processes that are ultimately responsible for kilonova signals. Both astrophysical scenarios involve strongly magnetized matter.

ACKNOWLEDGMENTS

This work was supported by the Max Planck Society's Independent Research Group Programme. J. V. K. kindly acknowledges the CARIPARO Foundation for funding his Ph.D. fellowship within the Ph.D. School in Physics at the University of Padova.

APPENDIX: DERIVATIONS

In this Appendix, we provide derivation steps left out in the main text. First, we derive Eq. (65) for the derivative of the master function. Starting from Eq. (44),

$$f = \mu - \hat{\mu} = \mu - (\hat{\nu} + \bar{r}^2\mu)^{-1}, \quad (\text{A1})$$

$$f' = 1 + (\hat{\nu} + \bar{r}^2\mu)^{-2}(\hat{\nu}' + \bar{r}^2 + 2\bar{r}\bar{r}'\mu) \quad (\text{A2})$$

$$= 1 + \hat{\mu}^2(\hat{\nu}' + \bar{r}^2 - 2(1-x)x^2r_\perp^2), \quad (\text{A3})$$

where primes denote derivatives with respect to μ . We first consider the case where ϵ computed from Eq. (42) does not exceed the upper limit allowed by the EOS (but may be smaller than the zero-temperature limit). For those cases, we can use Eq. (64) to get

$$f' = 1 + \hat{\mu}^2 \left(-(1+c_s^2)\hat{\nu} \frac{\hat{W}'}{\hat{W}} + \bar{r}^2 - 2(1-x)x^2r_\perp^2 \right) \quad (\text{A4})$$

$$= 1 + \hat{\mu}^2 \left(-(1+c_s^2)\hat{\nu}\hat{W}^2\mu(x^3r_\perp^2 + r_\parallel^2) + x^2r_\perp^2 + r_\parallel^2 - 2(1-x)x^2r_\perp^2 \right). \quad (\text{A5})$$

At a solution, we have $\hat{\mu} = \mu$ and $\hat{\nu}\hat{W}^2\mu = 1$, which leads to

$$f' = 1 + \mu^2 \left(-(1+c_s^2)(x^3r_\perp^2 + r_\parallel^2) + x^2r_\perp^2 + r_\parallel^2 - 2(1-x)x^2r_\perp^2 \right) \quad (\text{A6})$$

$$= 1 + \mu^2 \left((1-c_s^2)(x^3r_\perp^2 + r_\parallel^2) - x^2r_\perp^2 - r_\parallel^2 \right) \quad (\text{A7})$$

$$= 1 - \mu^2\bar{r}^2 + \mu^2(1-c_s^2)(x^3r_\perp^2 + r_\parallel^2). \quad (\text{A8})$$

Using the fact that $\hat{\nu} = \mu\bar{r}$ at the solution, we arrive at Eq. (65).

We now address the case where ϵ computed from Eq. (42) does exceed the limit ϵ_{\max} below which the EOS is valid. Equation (48) then becomes

$$\hat{\nu} = \nu_B = (1 + \hat{a})(1 + \bar{q} - \mu\bar{r}^2), \quad (\text{A9})$$

$$\hat{\nu}' = \frac{\hat{\nu}}{1 + \hat{a}}\hat{a}' + (1 + \hat{a})(\bar{q}' - \bar{r}^2 - 2\mu\bar{r}\bar{r}'). \quad (\text{A10})$$

Inserting Eqs. (56), (57), and (30) yields

$$\frac{\hat{\nu}'}{\hat{\nu}} = \frac{\hat{a}'}{1 + \hat{a}} - \frac{1 + \hat{a}}{\hat{\nu}}R^2, \quad (\text{A11})$$

with $R^2 \equiv x^3r_\perp^2 + r_\parallel^2 \leq \bar{r}^2$. For the case at hand, $\hat{a}(\hat{\rho}) = a(\hat{\rho}, \epsilon_{\max}(\hat{\rho}))$, and hence

$$\hat{a}' = \hat{\rho}' \left(\frac{\partial a}{\partial \rho} + \frac{d\epsilon_{\max}}{d\rho} \frac{\partial a}{\partial \epsilon} \right). \quad (\text{A12})$$

Splitting the derivative of ϵ_{\max} into adiabatic and residual contributions by using definition Eq. (67), we obtain

$$\hat{a}' = \hat{\rho}' \left(\frac{\partial a}{\partial \rho} + \left(\frac{\hat{P}}{\hat{\rho}^2} + \frac{A(\hat{\rho})}{\hat{\rho}} \right) \frac{\partial a}{\partial \epsilon} \right) \quad (\text{A13})$$

$$= \hat{\rho}' \left(\frac{da}{d\rho} \Big|_s + \frac{A(\hat{\rho})}{\hat{\rho}} \frac{\partial a}{\partial \epsilon} \right). \quad (\text{A14})$$

One can express the adiabatic sound speed in terms of a as

$$\frac{da}{d\rho} \Big|_s = \frac{1+a}{\rho}(c_s^2 - a), \quad (\text{A15})$$

which allows us to write

$$\hat{a}' = \frac{\hat{\rho}'}{\hat{\rho}} \left((1 + \hat{a})(c_s^2 - \hat{a}) + A \frac{\partial a}{\partial \epsilon} \right). \quad (\text{A16})$$

Evaluating at the solution, we can use $\hat{\mu} = \mu$ and $\hat{W}^2\hat{\nu}\mu = 1$. Using also Eqs. (59) and (60), we get

$$\frac{\hat{\nu}'}{\hat{\nu}} = -\hat{W}^2\mu R^2 \left(1 + c_s^2 + \frac{A}{1 + \hat{a}} \frac{\partial a}{\partial \epsilon} \right). \quad (\text{A17})$$

Inserting this into Eq. (A3), we can rewrite the master function derivative at the solution as

$$f' = 1 - \hat{\nu}^2 + \hat{\nu}^2 \frac{R^2}{\bar{r}^2} \left(1 - c_s^2 - \frac{A}{1 + \hat{a}} \frac{\partial a}{\partial \epsilon} \right). \quad (\text{A18})$$

If Eq. (66) holds, f' is always strictly positive at the solution, as claimed in Sec. II G.

Finally, we need to discuss the corner case where the mass density computed in Eq. (41) is outside the valid range of the EOS. If $\bar{D}/\hat{W}(\mu_+) > \rho_{\max}$ or $\bar{D} < \rho_{\min}$, there is no solution with valid density and no need to determine the root. The only complication arises when $\bar{D}/\hat{W}(\mu_+) < \rho_{\max} < \bar{D}$ or $\bar{D}/\hat{W}(\mu_+) < \rho_{\min} < \bar{D}$. Although the proof for the existence of a solution remains valid in those cases, we did not succeed in proving uniqueness of the master function root on the range $(0, \mu_+]$. Luckily, it is not necessary to prove uniqueness on the full interval. Instead, we numerically solve $\bar{D}/\hat{W}(\mu) = \rho_{\min/\max}$ for μ . We recall that $\hat{W}(\mu)$ is an analytic expression that does not

involve the EOS and monotonically increases with μ . Hence, we can find a subinterval $(\mu_a, \mu_b) \subseteq (0, \mu_+)$ that consists of all values for which the density $\bar{D}/\hat{W}(\mu)$ is valid. We already know that the master function f has at

most one root on that interval. There is a solution with valid density if and only if $f(\mu_a)$ and $f(\mu_b)$ have opposite signs, which is easy to check. Using an initial bracket (μ_a, μ_b) for the root finding then ensures uniqueness.

-
- [1] B. P. Abbott *et al.*, *Astrophys. J.* **848**, L12 (2017).
 [2] B. P. Abbott, R. Abbott *et al.* (Virgo and LIGO Scientific Collaborations), *Phys. Rev. Lett.* **119**, 161101 (2017).
 [3] B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. X* **9**, 011001 (2019).
 [4] B. P. Abbott *et al.*, *Astrophys. J.* **848**, L13 (2017).
 [5] E. Troja, L. Piro, H. van Eerten, R. T. Wollaeger, M. Im *et al.*, *Nature (London)* **551**, 71 (2017).
 [6] G. Hallinan, A. Corsi, K. P. Mooley, K. Hotokezaka, E. Nakar *et al.*, *Science* **358**, 1579 (2017).
 [7] J. D. Lyman, G. P. Lamb, A. J. Levan, I. Mandel, N. R. Tanvir *et al.*, *Nat. Astron.* **2**, 751 (2018).
 [8] D. Lazzati, R. Perna, B. J. Morsony, D. Lopez-Camara, M. Cantiello, R. Ciolfi, B. Giacomazzo, and J. C. Workman, *Phys. Rev. Lett.* **120**, 241103 (2018).
 [9] K. P. Mooley, A. T. Deller, O. Gottlieb, E. Nakar, G. Hallinan, S. Bourke, D. A. Frail, A. Horesh, A. Corsi, and K. Hotokezaka, *Nature (London)* **561**, 355 (2018).
 [10] G. Ghirlanda, O. S. Salafia, Z. Paragi, M. Giroletti, J. Yang *et al.*, *Science* **363**, 968 (2019).
 [11] R. Ciolfi, *Int. J. Mod. Phys. D* **27**, 1842004 (2018).
 [12] O. Just, M. Obergaulinger, H.-T. Janka, A. Bauswein, and N. Schwarz, *Astrophys. J. Lett.* **816**, L30 (2016).
 [13] A. Perego, H. Yasin, and A. Arcones, *J. Phys. G* **44**, 084007 (2017).
 [14] V. Paschalidis, M. Ruiz, and S. L. Shapiro, *Astrophys. J. Lett.* **806**, L14 (2015).
 [15] M. Ruiz, R. N. Lang, V. Paschalidis, and S. L. Shapiro, *Astrophys. J.* **824**, L6 (2016).
 [16] R. Ciolfi, *Mon. Not. R. Astron. Soc. Lett.* **495**, L66 (2020).
 [17] E. Pian, P. D’Avanzo, S. Benetti, M. Branchesi, E. Brocato *et al.*, *Nature (London)* **551**, 67 (2017).
 [18] D. M. Siegel and B. D. Metzger, *Astrophys. J.* **858**, 52 (2018).
 [19] R. Ciolfi and J. V. Kalinani, *Astrophys. J.* **900**, L35 (2020).
 [20] R. Ciolfi, *Gen. Relativ. Gravit.* **52**, 59 (2020).
 [21] A. J. Wright and I. Hawke, *Mon. Not. R. Astron. Soc.* **491**, 5510 (2020).
 [22] Z. B. Etienne, M.-B. Wan, M. C. Babiuc, S. T. McWilliams, and A. Choudhary, *Classical Quantum Gravity* **34**, 215001 (2017).
 [23] F. Galeazzi, W. Kastaun, L. Rezzolla, and J. A. Font, *Phys. Rev. D* **88**, 064009 (2013).
 [24] D. Radice and L. Rezzolla, *Astron. Astrophys.* **547**, A26 (2012).
 [25] B. Giacomazzo and L. Rezzolla, *Classical Quantum Gravity* **24**, S235 (2007).
 [26] P. Cerdá-Durán, J. A. Font, L. Antón, and E. Müller, *Astron. Astrophys.* **492**, 937 (2008).
 [27] P. Mösta, B. C. Mundim, J. A. Faber, R. Haas, S. C. Noble, T. Bode, F. Löffler, C. D. Ott, C. Reisswig, and E. Schnetter, *Classical Quantum Gravity* **31**, 015005 (2014).
 [28] Z. B. Etienne, V. Paschalidis, R. Haas, P. Mösta, and S. L. Shapiro, *Classical Quantum Gravity* **32**, 175009 (2015).
 [29] F. Cipolletta, J. V. Kalinani, B. Giacomazzo, and R. Ciolfi, *Classical Quantum Gravity* **37**, 135010 (2020).
 [30] S. C. Noble, C. F. Gammie, J. C. McKinney, and L. D. Zanna, *Astrophys. J.* **641**, 626 (2006).
 [31] D. Neilsen, S. L. Liebling, M. Anderson, L. Lehner, E. O’Connor, and C. Palenzuela, *Phys. Rev. D* **89**, 104029 (2014).
 [32] C. Palenzuela, S. L. Liebling, D. Neilsen, L. Lehner, O. L. Caballero, E. O’Connor, and M. Anderson, *Phys. Rev. D* **92**, 044045 (2015).
 [33] D. M. Siegel, P. Mösta, D. Desai, and S. Wu, *Astrophys. J.* **859**, 71 (2018).
 [34] W. Kastaun, RePrimAnd library V1.1, <https://doi.org/10.5281/zenodo.4075317> (2020).
 [35] J. S. Read, B. D. Lackey, B. J. Owen, and J. L. Friedman, *Phys. Rev. D* **79**, 124032 (2009).
 [36] H. Müller and B. D. Serot, *Nucl. Phys.* **A606**, 508 (1996).
 [37] R. Ciolfi, W. Kastaun, B. Giacomazzo, A. Endrizzi, D. M. Siegel, and R. Perna, *Phys. Rev. D* **95**, 063016 (2017).
 [38] R. Ciolfi, W. Kastaun, J. V. Kalinani, and B. Giacomazzo, *Phys. Rev. D* **100**, 023005 (2019).
 [39] W. Kastaun, R. Ciolfi, and B. Giacomazzo, *Phys. Rev. D* **94**, 044060 (2016).
 [40] A. Endrizzi, D. Logoteta, B. Giacomazzo, I. Bombaci, W. Kastaun, and R. Ciolfi, *Phys. Rev. D* **98**, 043015 (2018).
 [41] G. E. Alefeld, F. A. Potra, and Y. Shi, *ACM Trans. Math. Softw.* **21**, 327 (1995).
 [42] W. I. Newman and N. D. Hamlin, *SIAM J. Sci. Comput.* **36**, B661 (2014).