

Accurate evolutions of inspiralling neutron-star binaries: assessment of the truncation error

Luca Baiotti¹, Bruno Giacomazzo², and Luciano Rezzolla^{2,3}

¹ Graduate School of Arts and Sciences, University of Tokyo, Tokyo, Japan

² Max-Planck-Institut für Gravitationsphysik, Albert-Einstein-Institut, Potsdam, Germany

³ Department of Physics and Astronomy, Louisiana State University, Baton Rouge, LA, USA

Abstract. We have recently presented an investigation in full general relativity of the dynamics and gravitational-wave emission from binary neutron stars which inspiral and merge, producing a black hole surrounded by a torus [1]. We here discuss in more detail the convergence properties of the results presented in [1] and, in particular, the deterioration of the convergence rate at the merger and during the survival of the merged object, when strong shocks are formed and turbulence develops. We also show that physically reasonable and numerically convergent results obtained at low-resolution suffer however from large truncation errors and hence are of little physical use. We summarize our findings in an “*error budget*”, which includes the different sources of possible inaccuracies we have investigated and provides a first quantitative assessment of the precision in the modelling of compact fluid binaries.

PACS numbers: 04.30.Db, 04.70.Bw, 95.30.Lz, 97.60.Jd

1. Introduction

The study of the final stages of the evolution of binary systems is a cornerstone of any theory of gravity and a long-standing problem in General Relativity. Important issues in relativistic astrophysics still awaiting clarification, such as the mechanism responsible for gamma-ray bursts (GRBs), may be unveiled through a better understanding of the complex physics accompanying the inspiral and merger of two neutron stars. Furthermore, the study of the events that lead from a binary system of neutron stars to a black hole, possibly surrounded by a hot and high-density disc, will provide the gravitational waveforms and the energetics of one of the most important sources of gravitational radiation. While analytical techniques are very successful in describing binary systems which are widely separated and thus moving at comparatively small velocities and in slowly varying gravitational fields, numerical simulations represent possibly the only tool to investigate the dynamics when the two compact objects are performing the final few orbits of their evolution and the dynamics are in a fully nonlinear regime.

Considerable progress has been achieved in the past few years in the modelling in full general relativity of compact-objects binaries (see [2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12] for the most recent work from the different groups), and we have recently presented a systematic investigation of the dynamics and gravitational-wave emission from binary neutron stars which inspiral and merge, producing a black hole surrounded by a torus [1] (hereafter paper I). The purpose of this paper is to consider in more detail the results presented in paper I and assess critically their accuracy, their convergence properties and highlight the difficulties that can be encountered when simulating the turbulent motions that develop after the merger and

that characterize the subsequent evolution of the merged object up to the collapse to a black hole.

The plan of the paper is as follows. After a brief introduction in Section 2 to our code and to the binary models we simulated, we give in Section 3 a review of the general dynamics of the inspiral and merger of equal-mass neutron star binaries. Section 4 is instead dedicated to the discussion of the accuracy of our simulations when performed with different resolutions, while Section 5 assesses the influence that purely numerical aspects have on the physical results and presents them in a compact “*error budget*”. Finally, our summary and conclusions are collected in Section 6.

2. Mathematical and Numerical Setup

All the details on the mathematical and numerical setup used for producing the results presented here are discussed in depth in [13, 1]. In what follows, we limit ourselves to a brief overview.

2.1. Einstein and Hydrodynamics equations

The evolution of the spacetime was obtained using the `Ccatie` code, a three-dimensional finite-differencing code providing a solution of a conformal traceless formulation of the Einstein equations [13]. The general-relativistic equations were instead solved using the `Whisky` code presented in [14, 15, 16], which adopts a flux-conservative formulation of the equations as presented in [17] and high-resolution shock-capturing schemes. The `Whisky` code implements several reconstruction methods, such as Total-Variation-Diminishing (TVD) methods, Essentially-Non-Oscillatory (ENO) methods [18] and the Piecewise Parabolic Method (PPM) [19]. Also, a variety of approximate Riemann solvers can be used, starting from the Harten-Lax-van Leer-Einfeldt (HLLE) solver [20], over to the Roe solver [21] and the Marquina flux formula [22] (see [14, 15] for a more detailed discussion). All the results reported hereafter have been computed using the Marquina flux formula and a PPM reconstruction. We stress again (as already done in [1, 23]) that the use of high-order methods and high-resolution is *essential* to be able to draw robust conclusions on the inspiral and merger. Lower-order methods in the reconstruction and low resolution may yield convergent and apparently reasonable results which however contain a large truncation error. Specific examples of this type of problem are presented in Appendix 1 of paper I and in Figure 4 of [23].

The system of hydrodynamics equations is closed by an equation of state (EOS) and, as discussed in detail in [1], the choice of the EOS plays a fundamental role in the post-merger dynamics and significantly influences the survival time, against gravitational collapse, of the hyper-massive neutron star (HMNS) likely produced by the merger. It is therefore important that special attention is paid to use EOSs that are physically realistic, as done in [24] within a conformally flat description of the fields and a simplified treatment of the hydrodynamics. Because we are here mostly concerned with assessing the size of the truncation error rather than with a realistic description of the neutron-star matter, we have employed the commonly used “ideal-fluid” EOS, in which the pressure p is expressed as $p = \rho \epsilon (\Gamma - 1)$, where ρ is the rest-mass density, ϵ is the specific internal energy and Γ is the adiabatic exponent. Such an EOS, while simple, provides a reasonable approximation and we expect that the use of realistic EOSs would not change the main results of this work.

2.2. Adaptive Mesh Refinements

Both the Einstein and the hydrodynamics equations are solved using the vertex-centered adaptive mesh-refinement (AMR) approach provided by the `Carpet` driver [25]. Our rather basic form of AMR consists in centering the highest-resolution level around the peak in the rest-mass density of each star and in moving the “boxes” so to track the position of this maximum as the stars orbit. The boxes are merged when they overlap.

The results presented below refer to simulations performed at three different resolutions and for each of them we have used 6 levels of mesh refinement. More specifically, the finest refinement level has been chosen to have resolution of either $h = 0.1875 M_{\odot} = 277$ m (hereafter “low” resolution), of $h = 0.1500 M_{\odot} = 222$ m (hereafter “medium” resolution), or of $h = 0.1200 M_{\odot} = 177$ m (hereafter “high” resolution). It may be useful to point out that although the simulation at high-resolution stretches the computational resources available to us, it is effectively more expensive (both in memory and computing time) than an equivalent simulation carried out for binary black holes at resolutions that are 10 times larger but on fine grids that are much smaller (*cf.* the high-resolution in [13]). Using the medium resolution as a reference, the grid structure is such that the side of the finest grids is $15 M_{\odot} = 22.15$ km, while a single grid resolution covers the region between a distance $r = 150 M_{\odot} = 221.5$ km and $r = 250 M_{\odot} = 369.2$ km from the center of the domain. The latter region is the one in which our gravitational-wave extraction is carried out, with a resolution of $h = 4.8 M_{\odot} = 7.088$ km (as a comparison, the gravitational wavelength is ~ 100 km and thus reasonably well-resolved on this grid). In addition, a set of refined but fixed grids is set up at the center of the computational domain so as to capture the details of the Kelvin-Helmholtz instability (*cf.* paper I). Unless explicitly stated, for all the simulations reported here we have used a reflection-symmetry condition across the $z = 0$ plane and a π -symmetry condition across the $x = 0$ plane.

An important difference with respect to paper I, where the finest grid was covering only the central region of each neutron star, is that here each star is completely covered by the finest grid. Although this choice is computationally more expensive, it allows us to reach convergent results already with resolutions $h \gtrsim 0.19 M_{\odot}$, which are therefore smaller than those discussed in paper I (where we have used what is here the high resolution).

The timestep on each grid is set by the Courant condition (expressed in terms of the speed of light) and so by the spatial grid resolution for that level; the typical Courant coefficient is set to be 0.35. The time evolution is carried out using 4th-order-accurate Runge-Kutta integration algorithm. Boundary data for finer grids are calculated with spatial prolongation operators employing 3rd-order polynomials and with prolongation in time employing 2nd-order polynomials. The latter allows a significant memory saving, requiring only three timelevels to be stored, with little loss of accuracy due to the long dynamical timescale relative to the typical grid timestep. See Section 5 for a discussion on the changes caused by the different interpolation order.

2.3. Initial data

As initial data we use the general relativistic binaries produced by Taniguchi and Gourgoulhon [26] with the multidomain spectral-method code `Lorene` [27]. The initial solutions for the binaries are obtained assuming a quasi-circular orbit, an irrotational velocity field, and a conformally-flat spatial metric. The matter is modelled using a polytropic EOS $p = K \rho^{\Gamma}$ with $K = 123.6$ and $\Gamma = 2$.

In paper I, a number of different initial-data configurations for neutron-star binaries were

used in order to illustrate the variety of possible behaviours. Here, however, the measure of the truncation error is more easily done when considering a single configuration and in particular one that, although representative, also reduces the computational costs. As a result, we have chosen a model with a rather high mass so that a black hole is formed soon after the merger. More specifically, our fiducial binary has the following characteristics: a proper separation between the centers of the stars $d/M_{\text{ADM}} = 12.6 = 60.3$ km (this corresponds to a coordinate distance of ~ 45 km); a baryon mass of each star $M_b = 1.78 M_{\odot}$; a total ADM mass $M_{\text{ADM}} = 3.23 M_{\odot}$; an initial angular momentum $J = 10.13 M_{\odot}^2 = 8.92 \times 10^{49}$ g cm²/s; an initial orbital angular velocity $\Omega_0 = 9.39 \times 10^{-3} = 1.9$ rad/ms = 302 Hz; a ratio of the polar to the equatorial coordinate radius of each star $r_p/r_e = 0.945$.

A note of caution should be made. While our choice for a rather massive binary does have the advantage of allowing, within a total timescale of ~ 11 ms, for the analysis of *both* the inspiral (over more than three orbits) and of the merger, it also leads to a very rapid collapse to a black hole once the HMNS has been produced (*i.e.*, after only ~ 1 ms). This is simply because despite the high temperature and degree of differential rotation, the HMNS is so massive to be well beyond the instability threshold for the collapse to a black hole. As a result, this choice inevitably prevents us from determining how different choices of grids or resolutions influence the survival time of the HMNS when it is very close to the instability threshold. This is clearly a limitation of the present approach which will be overcome once larger computational facilities become available.

3. Review of the dynamics of the merger

In paper I we have described in detail the dynamics of the matter during the inspiral, the merger, the transition to the collapse, the collapse, and then the black-hole ringdown, for binaries with different initial masses, initial distances, and (idealized) EOSs. In what follows, we briefly summarize those results.

A first notable result of paper I is that of having clearly shown that for any given mass, the survival time of the HMNS depends on the EOS. More specifically, we have shown that the polytropic and isentropic EOS $p = K\rho^{\Gamma}$, for which therefore no shock heating is possible, leads either to the *prompt* formation of a rapidly rotating black hole surrounded by a dense torus in the higher-mass case, or, in the lower-mass case, to a HMNS, which develops a bar, emits large amounts of gravitational radiation and eventually experiences a *delayed* collapse to black hole. We have also shown that, for both initial masses, the ideal-fluid EOS inevitably leads to a further delay in the collapse to black hole as a result of the larger pressure support provided by the temperature increase due to shocks. In this case the temperature in the formed HMNS can reach values as high as $10^{11} - 10^{12}$ K, so that the subsequent dynamics and especially the time of the collapse can be reduced if cooling mechanisms, such as the direct URCA[‡] process, take place.

With the exception of the low-mass ideal-fluid binary, whose HMNS is expected to collapse to black hole on a timescale which is computationally challenging (*i.e.*, ~ 110 ms), all the binaries considered lead to the formation of a rotating black hole surrounded by a rapidly rotating torus. The masses and dimensions of the tori depend on the EOS, but are generically larger than those reported in previous independent studies, with masses up to $\approx 0.07 M_{\odot}$. Confirming what was reported in [28], we have found that the amount of angular momentum lost during the inspiral phase can influence the mass of the torus for binaries that

[‡] Direct URCA processes are processes that produce neutrinos, via the decay of a baryon or an interaction of a baryon with a lepton. The word URCA is not an acronym, but the name of the place where these processes were first discussed by G. Gamow and M. Schenberg.

have the same EOS. In particular, the binaries that lose less angular momentum during the inspiral, namely the comparatively *low-mass* binaries, are expected to have comparatively *high-mass* tori. In addition, we have also considered the comparison of binaries of the same mass but with different initial coordinate separation (*i.e.*, 45 and 60 km) and found that there is an excellent agreement in the inspiral phase (as expected from the lowest-order post-Newtonian approximations), but also small differences at the merger and in the subsequent evolution, most likely due to small differences in the initial data.

Besides the study of the large-scale dynamics of the two neutron stars, in paper I we have also investigated the small-scale hydrodynamics of the merger and the possibility that dynamical instabilities develop. In this way, we have provided the first quantitative description of the onset and development of the Kelvin-Helmholtz instability, which takes place during the first stages of the merger phase, when the outer layers of the stars come into contact and a shear interface forms. The instability curls the interface forming a series of vortices which we were able to resolve accurately using the higher resolutions provided by AMR techniques. This instability, which could have important consequences in the generation of large magnetic fields even from small initial ones, has been recently discussed in [23].

Special attention in paper I was also dedicated to the analysis of the waveforms produced and to their properties for the different configurations. In particular, we have found that the largest loss rates of energy and angular momentum via gravitational radiation develop at the time of the collapse to black hole and during the first stages of the subsequent ringdown. Nevertheless, the configurations which emit the largest amounts of energy and angular momentum are those with lower masses, since they do not collapse promptly to a black hole. Instead, they produce a violently oscillating HMNS, which emits copious gravitational radiation, while rearranging its angular-momentum distribution, until the onset of the collapse to black hole. We have also found that, although the gravitational-wave emission from neutron-star binaries has spectral distributions with large powers at high frequencies (*i.e.*, $\gtrsim 1$ kHz), a signal-to-noise ratio as large as 3 can be estimated for a source at 10 Mpc using the sensitivity of currently operating interferometric detectors.

4. Accuracy of the results

Having reviewed the general dynamics of the inspiral and merger, we next assess the truncation error of the results. As a first representative measure of the accuracy of our simulations, we report in the left panel of Figure 1 the evolution of the L_2 -norm of the Hamiltonian-constraint violation [*i.e.*, equation (10) in paper I] for simulations at high resolution (dashed line), at medium resolution (solid line) and at low resolution (long-dashed line). Although the resolutions are rather similar (their ratio is just 0.8), they are high enough to stretch the supercomputer facilities we have access to. Furthermore, as we will discuss below, they are only marginally sufficient to provide convergent results and resolutions lower than $h \simeq 0.19 M_\odot$ would lead to results that are only consistent \S . The curves relative to the high and low resolutions have been scaled by a suitable factor compensating for the different truncation error. The vertical lines are relative to the high-resolution run and indicate respectively the time of the merger t_{merg} (dotted line) and the time when an apparent horizon is first found t_{AH} (dot-dashed line). Taking the L_2 -norm of the same constraint as a measure of the average truncation error, Figure 1 shows that this is initially $\lesssim 10^{-6}$ and that it grows

\S We recall that a numerical solution is said to be *consistent* if the truncation error associated to it, ϵ_T , tends to zero in the limit of infinite resolution, *i.e.*, $|\epsilon_T| = 0$ for $h \rightarrow 0$. A numerical solution is said to be *convergent* at the order p if and only if $|\epsilon_T| = \kappa h^p$ for $h \rightarrow 0$, where κ is a positive constant and p is the *chosen* truncation order. Clearly, while a convergent solution is also consistent, the opposite is not necessarily true.

rapidly to $\lesssim 10^{-4}$ at the time of the merger and later when the black hole is formed. As a comparison, the L_∞ -norm of the violation of the Hamiltonian constraint grows from $\sim 10^{-3}$ initially to $\sim 10^{-2}$ just outside the apparent horizon; much larger violations, *i.e.*, $L_\infty \sim 1$, are present inside the apparent horizon and are typical of these calculations (*cf.* [29]).

Once the black hole has essentially stopped ringing, the violation does not increase further. The convergence rate measured before the merger, thus involving ~ 3 orbits, is therefore $\simeq 1.8$ as shown by the good overlap between the mid-resolution and the rescaled high/low-resolution curves. We note that: although *i)* the `Ccatie` code uses 4th-order spatial differential operators, *ii)* the time update is made with a 4th-order accurate Runge-Kutta integration scheme, and *iii)* the PPM reconstruction scheme is 3rd-order (at most), a convergence rate of $\simeq 1.8$ is not at all surprising and is indeed the same measured with PPM in much simpler configurations, such as in the evolution of isolated stars [15]||. This is simply due to the fact that the truncation error of HRSC schemes is not uniform across the computational domain and can drop to lower orders at local extrema or across shocks.

After the merger, the convergence rate drops to $\simeq 1.2$ and this deterioration may be due to at least two different reasons and is probably due to both. Firstly, we note that the merger is accompanied by very large and extended shocks and, as mentioned above, a general property of HRSC methods is to become only 1st-order accurate at discontinuities. Secondly, the merger leads to the development of a Kelvin-Helmholtz instability with the consequent generation of a turbulent regime (see the discussion in Section III E of paper I). Under these conditions, the whole notion of convergence needs to be revisited and a possible approach in these cases is the one presented in [30]. Note also that the convergence order increases again after the formation of the black hole, when only a small amount of matter is present in the form of an orbiting torus. Such flows do not have large-scale shocks, are much smoother and hence lead to convergence rates that are again ~ 1.8 . Clearly, tests at even higher resolutions should be performed to determine the precise cause of the degradation in the convergence order, but it is unlikely that these will be feasible with the present computational resources. As a final remark, it is worth reporting that a behaviour similar to the one discussed here has been found [31] also in the independent AMR code `SACRA` [5] (based on similar numerical methods).

The right panel of Figure 1 shows instead the evolution of the maximum of the rest-mass density when normalized to its initial value and where, for clarity, we have reported only the medium and high-resolution results (these will be the reference resolutions hereafter). Clearly, the two lines show a behaviour which is extremely similar during the inspiral phase, with only very small differences appearing at the time of the merger and after the collapse. These differences are caused mostly by the slight difference in the time of the merger, which is $\simeq 1.7\%$ smaller in the case of the lower-resolution simulation and which is probably due to the slightly different initial data and losses of angular momentum. Note that we have here used a technique, first reported in [29], where we do not excise any part of the computational domain, but we rather rely on suitable singularity-avoiding slicing conditions and on the addition of a small amount of numerical dissipation to avoid that small inaccuracies influence the portion of the spacetime exterior to the apparent horizon. Since we do not reserve a special treatment to the matter collapsing inside the apparent horizon, the matter variables and their spatial gradients become very large. These are reasonably well-handled by the HRSC schemes, but only up until all the matter is confined to a few grid cells; at that point,

|| When local maxima (*e.g.*, at a stellar centre) and/or small shocks in the vicinity of the stellar surface (see the discussion in Sect. III C of [1]) are present - that is, in the interval preceding the collapse in the simulations reported here - the convergence order is lower than third; fortunately these errors are small enough not to degrade the global convergence order to first order only.

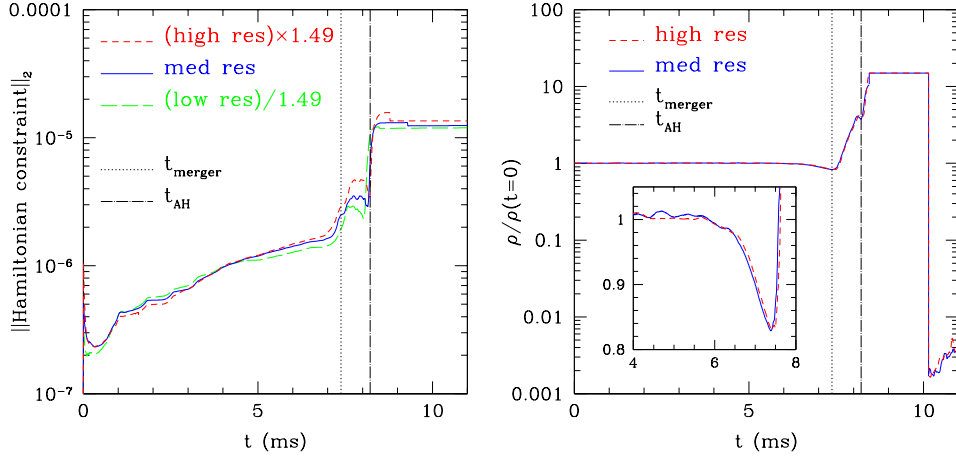


Figure 1. *Left panel:* Evolution of the L_2 -norm of the Hamiltonian-constraint violation for simulations at high resolution (dashed line), at medium resolution (solid line) and at low resolution (long-dashed line). Note that the high- and low-resolution curves are rescaled by 1.49, which corresponds to a convergence rate of 1.8. The vertical lines indicate the time of merger t_{merger} for the high-resolution run (dotted line) and the time when an apparent horizon is first found t_{AH} (dot-dashed line). *Right panel:* Evolution of the maximum of the rest-mass density normalized at the initial value and for the high and medium resolutions. Note that the horizontal lines mark the part of the evolution when the rest-mass density is unreliable after the formation of the apparent horizon; when the matter inside the apparent horizon has been dissipated, the maximum of the rest-mass density is that of the torus (see text for details).

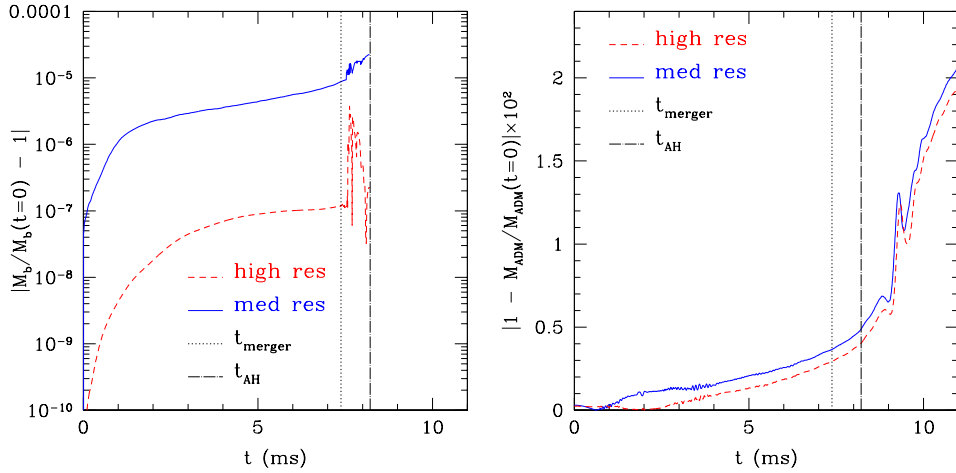


Figure 2. *Left panel:* Evolution of the error in the conservation of the rest mass M_b for the medium and high resolutions. Note that the fluctuations in M_b after the merger are due to the small amount of matter that ends up in buffer zones between two refinements levels. *Right panel:* Evolution of the normalized ADM mass as measured on the numerical grid; until when significant amounts of gravitational radiation leave the numerical grid (*i.e.*, up until about 9 ms), the differences between the two resolutions measure the error in computing M_{ADM} .

the conservation properties of our methods are grossly overstretched and the matter is simply dissipated. Because of these inaccuracies inside the apparent horizon, the evolution of the

maximum of the rest-mass density is unreliable after the formation of the apparent horizon and up until the rest-mass inside the apparent horizon has become negligible. This is shown with constant horizontal lines in the right panel of Figure 1. Note that when the matter inside the apparent horizon has been dissipated, *i.e.*, for $t \gtrsim 10$ ms, the maximum of the rest-mass density is located outside the apparent horizon and thus represents the maximum rest-mass density of the torus.

A first measure of the conservation properties of our simulations is offered in Figure 2, whose left and right panels show respectively the evolution of the error in the conservation of the rest mass M_b and of the ADM mass M_{ADM} . Note that thanks to the formulation of the hydrodynamics equations adopted and to the use of HRSC methods, the baryon mass is extremely well conserved, with deviations that are $\sim 10^{-6}$ ($\sim 10^{-7}$ for the high-resolution simulation) before the merger and are always below $\sim 10^{-5}$ ($\sim 10^{-6}$) up to the formation of the black hole. Note also that at the time of the merger the new grid structure leads to a certain amount of matter ending up in buffer zones between two refinements levels. This inevitably increases the error in the conservation, which remains anyway extremely good. This is particularly evident for the high-resolution run, where the error is intrinsically tiny. As discussed above, after the apparent horizon is found the baryon mass is no longer conserved up until all the matter inside the apparent horizon has been dissipated numerically. When this has happened, the baryon mass is the one contained in the torus and is equally well conserved, although it cannot be shown in the same panel because much smaller.

The ADM mass cannot be conserved when calculated on the finite-size numerical grid and it is rather expected to decrease as gravitational radiation leaves the numerical grid. The right panel of Figure 2 shows the evolution of the ADM mass of the two reference resolutions when normalized to their initial value. Although this quantity should not be interpreted as an error, the difference between the two resolutions gives a measure of the truncation error in computing M_{ADM} . Such a difference is $\sim 0.1\%$ for essentially all the simulations and the large variations at ~ 9 ms take place when the gravitational radiation produced by the collapse to black hole leaves the computational domain.

A second and more stringent measure of the overall conservation properties of our simulations is presented in Figure 3 and involves quantities which are partially radiated during the simulation. More specifically, the left panel shows the evolution of the total mass as normalized to the initial value and relative to the high-resolution simulation. Indicated with different lines are the volume-integrated values of the ADM mass (solid line), of the energy lost to gravitational waves (long-dashed line), and of their sum (dashed line). The last quantity should be strictly constant and this is the case to a precision of $\sim 0.5\%$ during the inspiral, but with a secular decreases that brings the total error to be $\sim 1\%$ at the end of the simulation (*cf.* dot-dashed line). Similar considerations apply also to the conservation of the angular momentum (*cf.* Figure 9 of paper I). This is shown in the right panel of Figure 3, which uses the same conventions as the left panel [here J_{vol} is computed with the integral (15) in paper I]. In this case the radiative losses are much larger (almost 15% of the available angular momentum is lost to gravitational waves) but, as shown in the figure, the overall conservation is accurate to $\sim 1\%$.

The accuracy of our gravitational-wave extraction is summarized in Figure 4, where in the left panel we show the gravitational-wave signal as computed via the $\ell = 2$, $m = 2$ component of the Weyl scalar quantity Ψ_4 , *i.e.*, $r(\Psi_4)_{22}/M_\odot$, as a function of the retarded time. The data is that of the high-resolution simulation and different lines refer to different extraction 2-spheres; the agreement among the three curves is extremely good, with the phases and amplitudes being almost coincident (see also the small inset). The right panel of Figure 4 shows instead $r(\Psi_4)_{22}/M_\odot$ at the high and medium resolutions, as extracted at $r = 200 M_\odot$.

As mentioned above, the slightly different initial data and losses of angular momentum will lead to a slight offset of the waveforms as the inspiral proceeds (this is of ~ 0.03 ms) and the data is shown once this offset has been removed, *i.e.*, after aligning the maxima of $|(\Psi_4)_{22}|$. Clearly the two waveforms are extremely similar and small differences can be seen only when looking at the total amplitude $r|\Psi_4|_{22}/M_\odot$ (see the inset).

In the left panel of Figure 5 we present a convergence test for the Ψ_4 waveform as extracted at $r = 200 M_\odot$, where different lines represent the differences between the waveforms suitably shifted in time and scaled to compensate for the different truncation errors. As discussed in detail in [13], the dominant source of error is a de-phasing which causes the lower resolution evolutions to “lag” behind the higher resolution. This delay is usually rather small and between $\sim 2 \times 10^{-2}$ and $\sim 5 \times 10^{-2}$ ms, but it is clearly visible when comparing the total amplitude of Ψ_4 as a function of time. If not properly taken into account, this error spoils the convergence tests, making the residuals appear as over or under convergent (*cf.* Figure 2 of [13]). This is obviously an artifact of the near cancellation of the lowest-order terms in the truncation error and it is induced by the small time-differences at different resolutions. We have removed this contamination by shifting the time coordinate of the low and medium resolution runs by the time interval needed to produce an alignment of the maxima of the emitted radiation (details on how to do this are discussed in Appendix A of [13]). Once this correction is made, the rescaled residual errors overlap reasonably well, indicating a convergence rate of 1.8; this is most likely due to the fact that the evolution is 2nd-convergent for most of the time and because the coefficients of the $\sim \mathcal{O}(h)$ truncation error are much smaller than the corresponding $\sim \mathcal{O}(h^2)$ one. To the best of our knowledge this is the first time that such a convergence test is shown on the gravitational waves from neutron-star binaries.

Finally in the right panel of Figure 5 we report a comparison of waveforms computed either enforcing a π -symmetry (blue continuous line) or not (red dashed line). Clearly the two are essentially indistinguishable on this scale and the small phase differences are shown

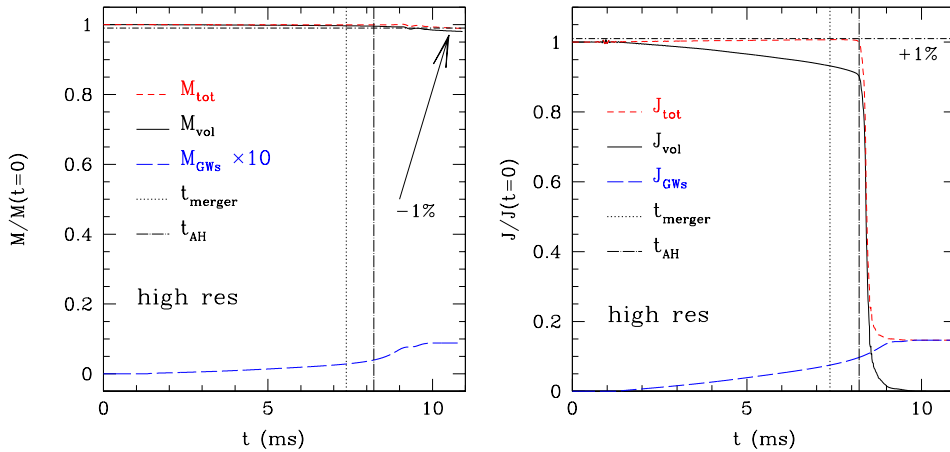


Figure 3. *Left panel:* conservation of energy. The continuous line is the ADM mass computed as an integral over the whole grid; The long-dashed line is the energy carried from gravitational waves outside the grid; The dashed line is the sum of the two and it should be conserved. The numerical violation is at most 1%. The data refer to the higher-resolution simulation. *Right panel:* the same as in the left panel but for the conservation of the angular momentum. Also in this case the violation is at most 1%.

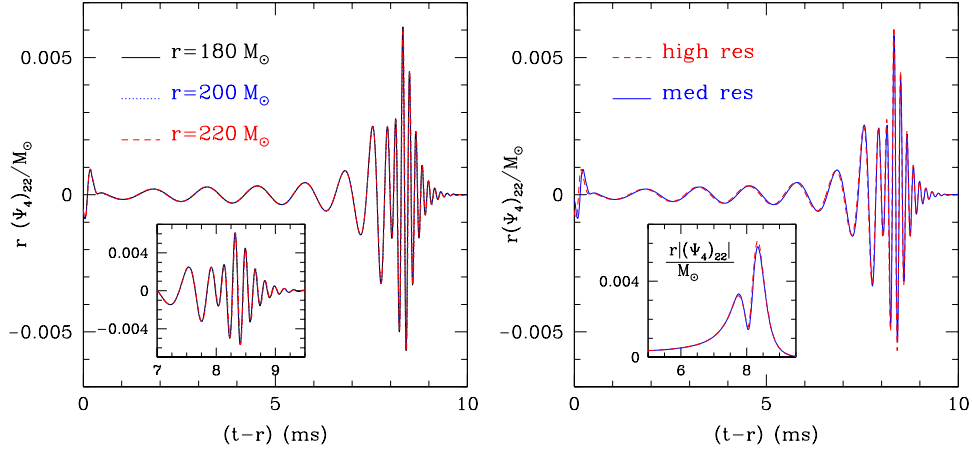


Figure 4. *Left panel:* $r(\Psi_4)_{22}/M_\odot$ waveform for different detectors, displayed versus retarded time. The data refer to the higher-resolution simulation. *Right panel:* $r(\Psi_4)_{22}/M_\odot$ waveform as computed for different resolutions.

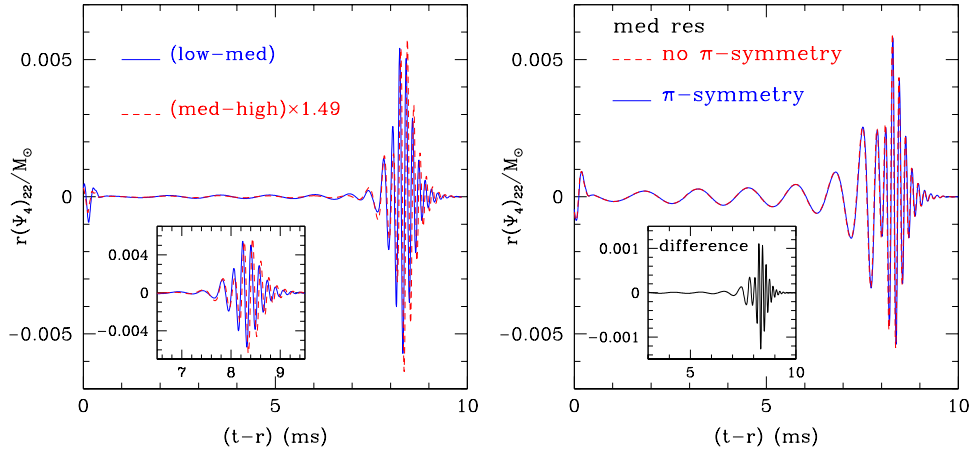


Figure 5. *Left panel:* Convergence test for $r(\Psi_4)_{22}/M_\odot$ waveform as extracted at $r = 200 M_\odot$. Different lines show the residuals between the waveforms computed at different resolutions, suitably shifted in time and scaled to compensate for a convergence rate of 1.8. *Right panel:* Waveforms computed either enforcing a π -symmetry (blue continuous line) or not (red dashed line). Indicated in the inset is the difference between the two waveforms.

in the inset.

5. Variations on the Theme and “Error Budget”

In the previous Sections we have discussed in detail the accuracy of our simulations for a fiducial high-mass binary by analysing how the results change with resolution. However, the resolution is not the only aspect of these simulations which can be modified to yield slightly different results, although it is unquestionably one of the most important ones. This Section is

Table 1. “Error-budget”, *i.e.*, variation of physical quantities induced by changes in the numerical settings. The table reports: the numerical setting varied, the amount of such a variation with respect to our standard simulation, and the percentile change of the time of the merger, $t_{\text{merg.}}$, and of that of the collapse, $t_{\text{coll.}}$ (see text for definitions), with respect to our standard simulation. A \pm sign indicates whether the time is increased (+) or decreased (−) with respect to the fiducial medium-resolution run.

Numerical setting varied	variation	$\Delta t_{\text{merg.}}$ (%)	$\Delta t_{\text{coll.}}$ (%)
resolution	$\times 1.25$ (from medium to high)	+0.45	+0.37
extent of finest grid	$\times 4$	+0.6	+0.6
spatial prolongation order	3rd vs. 5th order	+0.2	+0.3
artificial dissip. coeff.	$\times 100$ (from 0.1 to 0.001)	+0.1	+0.1
outer-boundary location	$\times 2$	−0.06	+0.1

therefore dedicated to those purely numerical aspects which can be varied and have been seen to yield (small) changes in the results. This assessment is obviously computationally very expensive, but it has allowed us to compile an “*error budget*” providing a simple reference on the influence that different sources of inaccuracies have on the physical results.

We build this budget by considering the evolution of the modulus of the waveform $|(\Psi_4)_{22}|$ (see, *e.g.*, Figure 20 of paper I or the inset in the right panel of Figure 4) and by singling out two specific times corresponding to the first and last maximum[¶]; we define these two times as the “time of merger”, $t_{\text{merg.}}$, and the “time of collapse”, $t_{\text{coll.}}$. For the high-resolution simulation described here, these are $t_{\text{merg.}} = 8.7$ ms and $t_{\text{coll.}} = 9.3$ ms, respectively. Note that the merger time defined above should not be confused with the time of the “hydrodynamical merger”, namely the time when the two stellar cores merge, $t_{\text{hd-merg.}}$, and which takes place considerably earlier ($t_{\text{hd-merg.}} = 7.4$ ms for the high-resolution case); we mark this hydrodynamical merger as the time when the maximum rest-mass density has a first significant minimum (*cf.* Figures 2 or 8 in [1]).

The error budget is reported in Table 1, whose rows contain information about the specific numerical setup which has been varied, the amount of the variation and the differences in the two times. Note that because the HMNS collapses promptly to a black hole, the differences in $t_{\text{coll.}}$ reported here are particularly small and should therefore be taken as lower limits. Indeed, for less massive HMNSs, the corresponding differences can be as large as $\sim 10\%$ and increase to $\gtrsim 70\%$ when considering simulations run at low resolution. The determination of $t_{\text{coll.}}$ represents therefore an emblematic example of how physically reasonable and numerically convergent results obtained at low-resolution suffer from very large truncation errors.

The first row in Table 1 reports the variations due to an increase in the resolution (*i.e.*, from the medium resolution to the high one) and does not need further comment besides noting that the results for the subsequent rows refer to medium-resolution simulations. The second row shows the change induced when the extent of the finest grid is increased by a factor of 4 in all the three spatial directions. In this case, a single fine grid covers both stars, so that the volume of the finest grid is 4^3 times larger and the simulation is therefore about 64 times more expensive. The third row reports the variations measured when the interpolation for the spatial prolongation operation needed in the time-stepping of our mesh refinements is changed from 3rd to 5th-order (only for the spacetime variables). The fourth row refers to a variation in the coefficient of the artificial-dissipation (see discussion in [29, 16]). The

[¶] When the collapse to black hole is delayed, the HMNS emits large amounts of gravitational radiation, thus leading to a series of local maxima in the evolution of $|(\Psi_4)_{22}|$ (see, *e.g.*, Figure 26 of paper I)

value of the dissipation coefficient ε in the fiducial simulation is 0.1 and the table reports the changes when using instead $\varepsilon = 0.001$; clearly the differences are minute but a higher dissipation also provides waveforms which are less affected by reflections among different refinement levels. Finally, the fifth row reports the changes when the outer-boundary location is increased by a factor 2. While the content of the table is self-explanatory, it is worth remarking that in all cases the induced variations are below 1% at these resolutions and that the times are essentially always increased.

6. Conclusions

We have presented a detailed analysis of the accuracy and convergence properties of our general-relativistic simulations of the inspiral and merger of binary neutron stars. More specifically, we have shown that for a high-mass binary and typical resolutions of $h \sim 0.19 M_\odot - 0.12 M_\odot$, the results show the expected convergence rate of 1.8 during the inspiral phase, the collapse and the subsequent ringdown. However, the convergence rate drops to 1.2 at the merger and during the evolution of the HMNS. This deterioration of the convergence rate is not surprising and can be explained as a combination of (at least) two effects. Firstly, strong shocks form during the merger and HRSC schemes are only 1st-order accurate at discontinuities. Secondly, the merger gives rise to turbulent motions, related to the development of a Kelvin-Helmholtz instability, for which the concept of local convergence needs to be revisited.

We have also presented tests on the conservation of physical quantities, such as the rest mass, the energy and the angular momentum, taking into account that the last two are not conserved because of losses in gravitational-wave. In particular, we have shown that, at our best resolution, the rest-mass is conserved with an error $\lesssim 10^{-6}$, while the energy and the angular momentum are conserved to $\lesssim 1\%$ after taking into account the parts lost to radiation. In addition, when considering the accuracy of the extraction of gravitational-wave information, we have shown the very good agreement, in both phase and amplitude, of the gravitational waves extracted from different detectors within the same simulation, or from the same detector but at different resolutions. Such waveforms have been shown to be also convergent at a rate of 1.8. Finally, we have reported a first investigation on how purely numerical changes in some of the setting of the simulation can influence the physical results. This analysis has led to the construction of a straightforward “error budget”, whose entries are all below 1%, at least for the high-mass binary considered here.

A final remark should be made when contrasting binary-black-hole simulations with the corresponding ones involving binary neutron stars. While for the first ones the numerical methods are sufficiently robust (and the convergence order sufficiently high) that an increase in resolution is usually the solution to the most serious problems, the hydrodynamical complexities inherent to the second ones (and the small convergence order after the merger) are such to require extra caution and very careful assessment of the possible sources of error. Experience has shown us that results at low resolution which appear reasonable and convergent, are however contaminated by large truncation errors and hence of little physical relevance. In view of this, and of the low convergence rate after the merger, we conclude with a remark that obvious in general, but is worth making when considering binary neutron star calculations: high-order methods and the highest possible resolutions are imperative to draw robust conclusions.

Acknowledgments

We thank the developers of Lorene for providing us with initial data, those of Carpet for the mesh refinement, and those of Cactus for the computational infrastructure our code is based on. Useful comments from Carlos Palenzuela and Masaru Shibata are kindly acknowledged. The computations were performed at the AEI and at LONI (www.loni.org). This work is also supported by the DFG SFB/Transregio 7 and by the JSPS grant 19-07803.

References

- [1] Baiotti L, Giacomazzo B and Rezzolla L 2008 *Phys. Rev. D* **78** 084033
- [2] Anderson M, Hirschmann E W, Lehner L, Liebling S L, Motl P M, Neilsen D, Palenzuela C and Tohline J E 2008 *Phys. Rev.* **D77** 024006
- [3] Anderson M, Hirschmann E W, Lehner L, Liebling S L, Motl P M, Neilsen D, Palenzuela C and Tohline J E 2008 *Phys. Rev. Lett.* **100** 191101
- [4] Duez M D *et al.* 2008 *Phys. Rev.* **D78** 104015
- [5] Yamamoto T, Shibata M and Taniguchi K 2008 *Phys. Rev. D* **78** 064054–1–38
- [6] Liu Y T, Shapiro S L, Etienne Z B and Taniguchi K 2008 *Phys. Rev. D* **78** 024012
- [7] Etienne Z B, Liu Y T, Shapiro S L and Baumgarte T W 2009 *Phys. Rev. D* **79** 044024
- [8] Rezzolla L *et al.* 2008 *Astrophys. J.* **674** L29–L32
- [9] Baker J G *et al.* 2008 *Phys. Rev.* **D78** 044046
- [10] Campanelli M, Lousto C O, Nakano H and Zlochower Y 2008 *arXiv:0808.0713*
- [11] Gonzalez J A, Sperhake U and Bruegmann B 2008 *arXiv:0811.3952*
- [12] Scheel M A, Boyle M, Chu T, Kidder L E, Matthews K D and Pfeiffer H P 2009 *Phys. Rev. D* **79** 024003
- [13] Pollney D, Reisswig C, Rezzolla L, Szilágyi B, Ansorg M, Deris B, Diener P, Dorband E N, Koppitz M, Nagar A and Schnetter E 2007 *Phys. Rev.* **D76** 124002
- [14] Baiotti L, Hawke I, Montero P and Rezzolla L 2003 *Computational Astrophysics in Italy: Methods and Tools* vol 1 ed Capuzzo-Dolcetta R (Trieste: MSAIS) p 210
- [15] Baiotti L, Hawke I, Montero P J, Löffler F, Rezzolla L, Stergioulas N, Font J A and Seidel E 2005 *Phys. Rev. D* **71** 024035
- [16] Baiotti L, Hawke I and Rezzolla L 2007 *Class. Quantum Grav.* **24** S187–S206
- [17] Banyuls F, Font J A, Ibáñez J M, Martí J M and Miralles J A 1997 *Astrophys. J.* **476** 221
- [18] Harten A, Engquist B, Osher S and Chakrabarty S R 1987 *J. Comput. Phys.* **71** 231
- [19] Colella P and Woodward P R 1984 *J. Comput. Phys.* **54** 174
- [20] Harten A, Lax P D and van Leer B 1983 *SIAM Rev.* **25** 35
- [21] Roe P L 1981 *J. Comput. Phys.* **43** 357
- [22] Aloy M A, Ibáñez J M, Martí J M and Müller E 1999 *Astrophys. J. Supp.* **122** 151
- [23] Giacomazzo B, Rezzolla L and Baiotti L 2009 *arXiv:0901.2722*
- [24] Oechslin R and Janka H T 2007 *Phys. Rev. Lett.* **99** 121102
- [25] Schnetter E, Hawley S H and Hawke I 2004 *Class. Quantum Grav.* **21** 1465–1488
- [26] Taniguchi K and Gourgoulhon E 2002 *Phys. Rev. D* **66** 104019
- [27] URL <http://www.lorene.obspm.fr>
- [28] Shibata M and Taniguchi K 2006 *Phys. Rev. D* **73** 064027
- [29] Baiotti L and Rezzolla L 2006 *Phys. Rev. Lett.* **97** 141101
- [30] Fromang S and Papaloizou J 2007 *Astron. Astrophys.* **476** 1113–1122
- [31] Shibata M, private communication, 2009