Computational Relativistic Astrophysics With Adaptive Mesh Refinement: Testbeds

Edwin Evans[1], Sai Iyer[1], Erik Schnetter[2], Wai-Mo Suen[1,3],
Jian Tao[1], Randy Wolfmeyer[1], and Hui-Min Zhang[1]

[1] McDonnell Center for the Space Sciences, Department of Physics,
Washington University, St. Louis, Missouri 63130
[2] Albert-Einstein-Institut, Am Mühlenberg 1, D-14467, Potsdam, Germany and
[3] Physics Department, Chinese University of Hong Kong, Hong Kong

We have carried out numerical simulations of strongly gravitating systems based on the Einstein equations coupled to the relativistic hydrodynamic equations using adaptive mesh refinement (AMR) techniques. AMR coalescences of neutron stars can now be simulated with sufficient resolution covering the neutron stars while having the computational domain extend to the local wave zone. We show an AMR simulation carried out with a workstation having an accuracy equivalent to that of a 1025^3 regular unigrid simulation, which is, to the best of our knowledge, larger than all previous simulations of similar NS systems on supercomputers. We believe the capability opens new possibilities in general relativistic simulations.

a. Introduction Numerical study of compact systems has received much attention due to observations in high-energy astronomy and the promise of gravitational wave astronomy. Most effort focuses on solving the Einstein equations with finite differencing methods. The main difficulty of this approach is that many general relativistic astrophysical processes of interest, e.g., processes involving black holes and neutron stars, require computational resources that are beyond what present day computers allow. The reasons that they are computationally demanding are 1. the lack of symmetry in realistic astrophysical situations, requiring the solving of the full set of Einstein equations coupled to the general relativistic hydrodynamic (GRHydro) equations in full 3+1 dimensional spacetime; and 2. the involvement of many length scales.

The difficulty of multiple length scales can be illustrated with the neutron star (NS) coalescence problem, one of our main systems of study. The length scales involved are: (i) A short length scale coming from the internal dynamics of a neutron star as a self gravitating object. One needs to be able to resolve the density variation to a reasonable accuracy before one can maintain a stable configuration in the Einstein theory. (ii) A longer length scale coming from dynamics of two neutron stars moving under the influence of one another, i.e., the length scale of the orbital separation between the two stars. (iii) The dynamical time scale of the system (the orbital period T) turns into a long length scale due to the dynamical nature of Einstein gravity (no such difficulty exists in Newtonian gravity, where one can evolve orbiting system more easily). The space surrounding the NSs within the corresponding length scale (the wavelength of the gravitational wave due to the orbital motion) needs to be covered in the computational domain as part of the system, both for the extraction of the waveform *and* for an accurate dynamical evolution (the problem manifests itself in the dynamical evolution being affected by the position of the outer boundary if put too close). (iv) The secular evolution time scale of the orbital motion turns into a resolution requirement for the numerical simulation as computational error accumulates. Our study in a previous paper [1] indicates that

1. To simulate a single isolated NS in a stable fashion with the Einstein equations requires a resolution on the order of $0.1M_0$, where M_0 is the baryonic mass of the neutron star, for a typical equation of state (EOS).

2. To set initial data in a fashion (e.g., using the conformally flat quasi-equilibrium (CFQE) approach) that we can have some confidence of its being astrophysically relevant, the initial separation of the two NSs would have to be on the order of $50M_0$ (depending on the initial spin states of the two NSs).

3. To get inspiral dynamics without much artificial influence from the boundary of the computational domain, the outer boundary has to be put at at least 0.5λ away, where λ is the gravitational wavelength of the system (assuming the present state of the art in setting outer boundary conditions for the constrained system of the Einstein equations).

4. To be able to accurately extract a gravitational waveform from the simulation, the computational domain should include up to 1λ.

5. To be able to evolve the spiraling NSs within the convergence regime to the point of coalescence: This depends on the choice of initial configuration and the numerical method used. With all existing methods we know of, the longer in time one needs to stay within the convergence regime (i.e., the constraint violations converging with respect to increasing resolution throughout that time period), the finer the resolution has to be. In our simulation reported in [1], a 643^3 simulation with $\Delta x = 0.2M_0$, covering up to 0.28λ for orbiting NSs at an initial separation of $28M_0$ (with an angular frequency of $\Omega = 0.012M_0^{-1}$), the system remains in the convergence regime for only about half an orbit. Being so far off

from our target of evolving to the coalescence point, an estimate of what might be needed would be meaningless.

The wavelength of a gravitational wave with orbital separation of $50M_0$ (cf., (2) above) is about $1,000M_0$. For a unigrid at $0.2M_0$ (cf., (1) above), the requirements (3) and (4) imply a grid of $5,000^3$. A computer with a memory size capable of doing such a simulation will not be available in the near future.

Due to this the biggest obstacle we encounter in NS coalescence simulation based on finite differencing of the Einstein equations is the need for a large number of grid points, which translates into large computer memory and long execution time. We need the adaptive mesh refinement (AMR) treatment: Use fine grid patches co-moving with the compact objects to satisfy the resolution required by (1), and a coarse grid extending to the local wave zone for (3) and (4). Similar considerations have motivated much effort in this direction, see e.g., [2, 3, 4, 5, 6, 7] for recent progress.

Unfortunately, application of AMR techniques in general relativistic astrophysics is more difficult than one might naively think. Although the theory and algorithms of mesh refinement are well established in computational science, and the numerical treatments of the Einstein equations and relativistic hydrodynamic equations have also been extensively investigated by relativists and astrophysicists, after many years of intense effort by many research groups it has not been possible to put the two together for a fully general relativistic 3D AMR simulation. The main difficulty is that it involves a huge infrastructure on *both* the computer science side and the physics side: it is difficult for computer scientists to dive into the complexity of the physics, and vise versa. As a rough representation of the complexity, in our code construction process, we have to integrate a 100,000 line mesh refinement code (GrACE [8]), a 85,000 line general relativistic astrophysics code (GR-Astro [9]) and a 500,000 line parallel computational library (Cactus Toolkit [10]) that GRAstro makes use of.

In this paper we demonstrate for the first time that a full 3+1 dimensional simulation based on the Einstein equations can be carried out with AMR. Three sample systems are studied:

1. A NS moving at a speed of $0.5c$ described by the general relativistic hydrodynamic equations coupled to the Einstein equations. The validity of our AMR treatment is examined with convergence tests. Convergence tests are more complicated with AMR; three different kinds of convergence tests are presented: (i) simulations with increasing resolutions on all grid levels, (ii) simulations with added levels of refinement, and (iii) comparison to unigrid results. The investigation of a boosted star, which invokes *all* terms in the evolution equations, played an important role in our code construction process.

2. Two NSs coalescing with angular momentum ($L = 8M_\odot{}^2$). The focus is on demonstrating that the AMR simulation can handle NS coalescences with an accuracy

comparable to that of an unigrid run with resolution same as the finest resolution of the AMR run.

3. An inspiraling neutron star binary. The two NSs are covered by co-moving fine grid patches, with the coarse grid covering a fraction of a wavelength of the system. We demonstrate an AMR simulation which is equivalent to a regular 1025^3 unigrid simulation, larger than any simulation of NS binary systems performed so far.

In the next 3 sections we discuss these three simulations. The last section summarizes and discuss the next steps.

b. Boosted Neutron Star. We begin with a study of a NS moving across an otherwise empty space at a constant speed. Although the physical system is not changing in time beyond a uniform boost, the metric has complicated spacetime dependences due to the frame dragging effect. Accordingly, all coordinate quantities including those of the spacetime and matter are changing in a non-trivial manner (not just a uniform translation). In the simulation, we start with a configuration satisfying the Hamiltonian and momentum constraint equations representing a NS boosted to $0.5c$, and evolve it with the full set of dynamical Einstein equations coupled to the GRHydro equations. The system of equations as well as the conventions we use in this paper are given in [1]. The simulation provides a good test for our code as it invokes *all* terms in the equations, and is numerically a fully dynamical test.

The NS is described by a polytropic equation of state (EOS): $P = (\Gamma - 1)\rho\epsilon = k\rho^\Gamma$, with $\Gamma = 2$ and $k = 0.0445c^2/\rho_n$, where ρ_n is the nuclear density (approximately 2.3×10^{14} g/cm^3). The NS has a proper radius of $R = 12M_\odot$, an ADM mass of $1.4M_\odot$ and a baryonic mass $M_0 = \frac{1}{2}\int d^3x \sqrt{\gamma}\rho W = 1.49M_\odot$. (For these values of parameters, the maximum stable NS configuration has an ADM mass of $1.79M_\odot$ and a baryonic mass of $1.97M_\odot$). The initial data is obtained by imposing a boost on the TOV solution ([11]). The evolution is carried out with the Γ freezing shift and the "$1 + \log$" lapse (for details of the shift and lapse conditions and method of implementations, see [1]).

The computational grid is set up as follows: 1. The coarse grid has a resolution of $dx = 2.88M_\odot$ (4 points across the radius R of the star) covering a region of $58M_\odot \times 58M_\odot \times 58M_\odot$. 2. Two levels of adaptive fine grid with $dx = 1.44M_\odot$ and $dx = 0.72M_\odot$ are set up. The adaptive grid is allowed to change in size and location as the refinement criteria dictate. 3. Two different refinement criteria have been studied: (i) value of matter density ρ, and (ii) amount of Hamiltonian constraint (HC) violation. Combinations of the two with "or" can be used. It turns out that for the neutron star studies it does not matter much which condition is used: the central region of the NS is at the same time the region of highest density, maximum HC violation and maximum evolution error. All simulations shown in this paper are obtained with (i).

In panel 1 of fig. 1 we show density profiles (represented as a height field) on the equatorial plane of the NS at 3
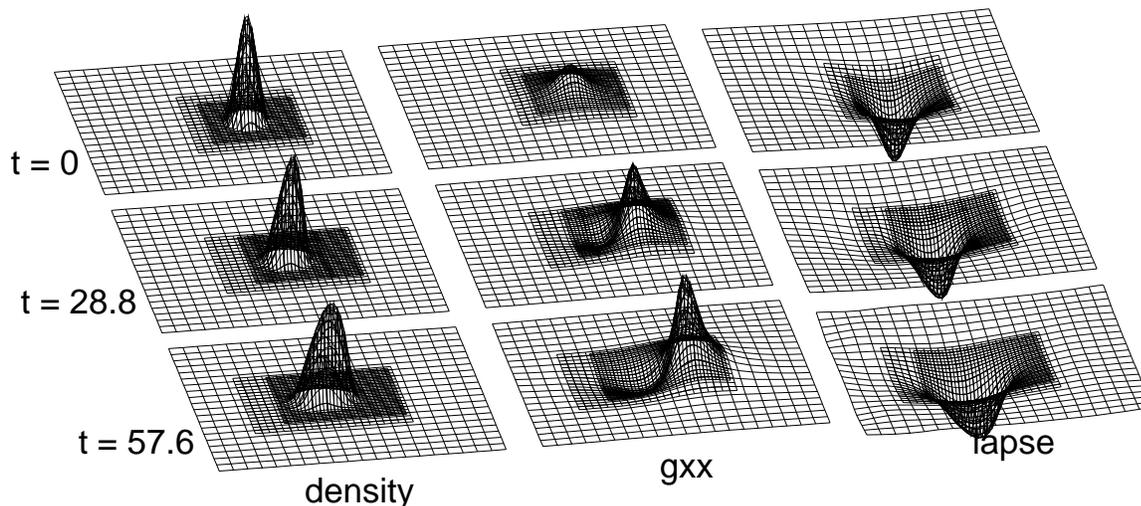
FIG. 1: Density, g_{xx}, and lapse of a boosted star on 3 grid levels are shown with 1 to 2 downsampling (only every other computational grid point shown), showing the motion of the star through the computational grid with distortion due to the frame dragging effect.
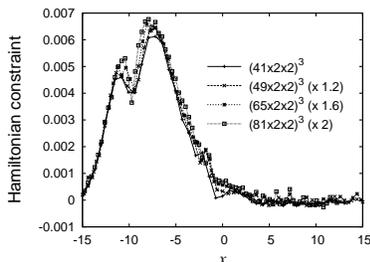
FIG. 2: HC violation scaled by $(1/dx)$ for runs with different resolutions. Their overlapping with one another implies first order convergence (as the TVD hydro scheme dictates).
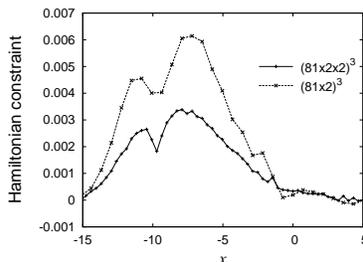
FIG. 3: Effect of more levels: the addition of one refinement level lowers the HC violation by a factor of 2 in the region of extra grid level where the HC violation is significant.
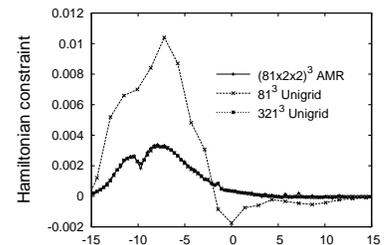
FIG. 4: The HC violation of the unigrid run (321^3) with resolution the same as that of the finest grid of the AMR run is nearly exactly the same as that of the AMR run. The unigrid run (81^3) with resolution the same as that of the coarse grid of the AMR run has much larger HC violation.

different times: $t = 0, 28.8, 57.6 M_\odot$. The grid structure is superimposed with a down sampling ratio of 2 (showing only every other grid point). We see that the NS as well as the fine grids are moving across the computational domain. In panel 2 we show the corresponding metric component g_{xx} with x being the direction of the motion. The frame dragging effect is readily seen. Panel 3 shows the evolution of the lapse.

Fig. 2–4 examine the validity of the simulation with 3 kinds of convergence tests. Fig. 2 shows the violation of the Hamiltonian constraint (HC) at $t = 28.8 M_\odot$ along the x-axis for four different runs. The HC violation is calculated on the finest grid available for regions covered by more than one grid. (This applies to all HC plots in this and the following sections.) The resolutions of the runs are 41^3, 49^3, 65^3, and 81^3, corresponding to $dx = 2.88 M_\odot, 2.4 M_\odot, 1.8 M_\odot, and 1.44 M_\odot$, respectively, on the base grid. (The notation $(41 \times 2 \times 2)^3$ indicates a 41^3 base grid and two levels of refinement with a refinement ratio

of 2 each.) The results for the higher resolution runs have been scaled linearly. The plot demonstrates that the code is converging to first order, which is the expected rate of convergence as we used a high resolution shock capturing TVD scheme [1] in our hydrodynamic evolution which is first order at extremal points.

In fig. 3, we compare the HC violations of two runs at time $t = 28.8 M_\odot$: (i) the $(81 \times 2 \times 2)^3$ run shown in fig. 2, and (ii) an $(81 \times 2)^3$ run with only one level of refinement covering the high density region. We see that the addition of a refinement level lowers the HC violation by a factor of 2 in the region of the extra grid level (where the HC violation is significant).

In fig. 4, we show the HC violations of three runs at $t = 28.8 M_\odot$. The AMR run is again the $(81 \times 2 \times 2)^3$ one in fig. 2. The other two are unigrid runs, one at the resolution of the coarsest AMR grid ($dx = 1.44 M_\odot$), and the other at the resolution of the finest AMR grid ($dx = 0.36 M_\odot$). We see that the AMR run has exactly

(a) $t = 0$ (b) $t = 56.6 M_\odot$ (c) $t = 96.4 M_\odot$
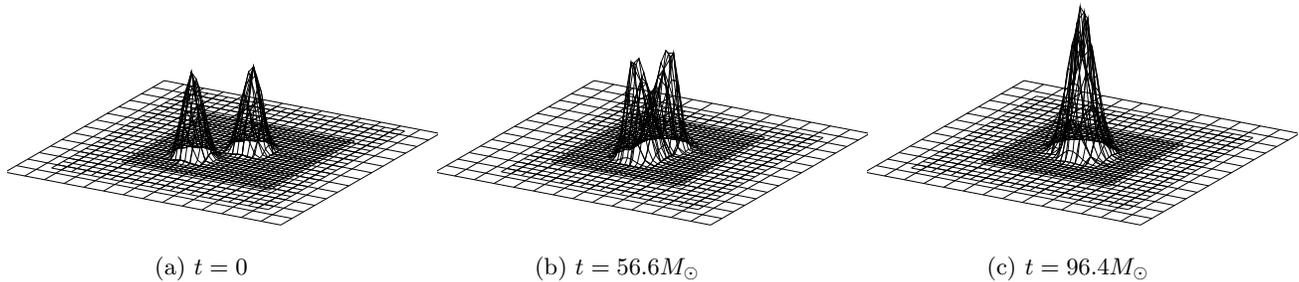
FIG. 5: Density profiles for coalescing stars at 3 different times, with the grid structure superimposed, and with 1 to 2 downsampling.

the same accuracy as the unigrid fine resolution run (the two lines coincide). This is an important point for our study: For NS simulations in this and the following sections, finite difference error is most significant in the high density region covered by the finest AMR grid; using coarser grids elsewhere does not affect the accuracy of the simulation. This enables us to speak of the "unigrid equivalent" of an AMR run: a unigrid run with the resolution of the finest AMR grid.

The three kinds of convergence tests provide confidence in the validity of our AMR treatment.

c. Coalescing Neutron Stars. In this section we study the coalescence of two NS's each having a baryonic mass and EOS as given above. The NSs have their equatorial plane on the x-y plane and an initial center to center (points of maximum mass) separation of $2.2R$ both in the x and y directions ($\sim 3R$ separation diagonally, $R = 12M_\odot$). Initially one NS is moving in the x direction and the other in the $-x$ direction with the total angular momentum of the system equal to $3.8M_0{}^2 = 8.4M_\odot{}^2$, which is roughly the angular momentum of a co-rotational conformally flat quasi-equilibrium configuration on the inner-most-stable circular orbit [1].

To determine the metric we have to solve the Hamiltonian and momentum constraint equations. The elliptic system was solved on a unigrid of 256^3 at resolution of $dx = 0.3M_\odot$. The initial metric and hydrodynamic data were then interpolated onto the AMR grids. The evolution obtained with the AMR treatment is then compared with that of the unigrid. The Γ freezing shift and $1 + \log$ slicing condition are used in both runs.

In fig. 5 we show density profiles (represented as height field) on the equatorial plane of the NSs at 3 different times $t = 0, 56.6, 94.4M_\odot$ in the AMR simulation, with the grid structure superimposed (downsampled by a factor of 2). At $t = 56.6M_\odot$, the two NSs start merging, and at $t = 94.4M_\odot$ the peaks of the density merged into one (at even later time two peaks reappear, having large angular momentum in the system).

In figs. 6a, 6b we compare the AMR results with the 257^3 unigrid run that has a resolution the same as the finest AMR grid. The lapse and g_{xx} along the x-axis
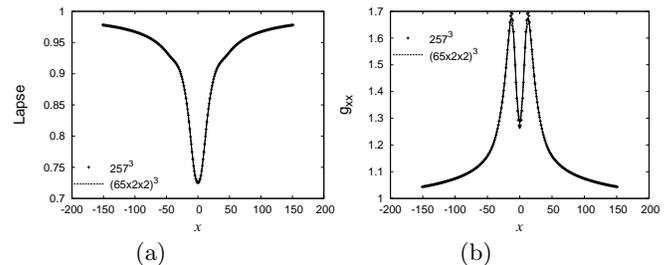
FIG. 6: Comparison of AMR (solid line) and unigrid (daggers) for coalescing neutron stars: (a) the lapse, (b) g_{xx}. Both agree to high accuracy.

at $t = 28M_\odot$ are given respectively in figs. 6a and 6b, with the AMR results given by the solid lines and the unigrid results by daggers. g_{xx} is shown as it is the metric function which shows the biggest difference between the AMR and unigrid runs. We see that the two runs give basically the same results. In fig. 7 we plot the values of the (spatial) maximum of the HC violations (which is one of the most sensitive measure of differences between runs) over time for the unigrid and AMR simulations. 16π times the maximum of ρ is also shown for comparison ($16\pi\rho$ is the source term in the HC equation). We see that at earlier time the error of the AMR run is slightly larger, while the error in the unigrid run picks up faster at later time, becoming larger than $16\pi\rho$. (At that point the simulation is no longer reliable as in some region the dynamical evolution is driven more by the HC violation than the matter density.)

d. Inspiraling Neutron Stars. In this section we demonstrate that with AMR we can now carry out on a workstation (Dell entry level server Poweredge 1850) NS inspiral simulations that is beyond existing unigrid simulations on supercomputers.

The NSs are taken to be initially in a conformally flat quasi-equilibrium (CFQE) irrotational circular orbit with an orbital separation of $3.3R$. Each NS has a baryonic mass of $1.625M_\odot$ with the same polytropic EOS as in the previous section. The initial data is obtained by solving the CFQE equations using the pseudo-spectral code de-
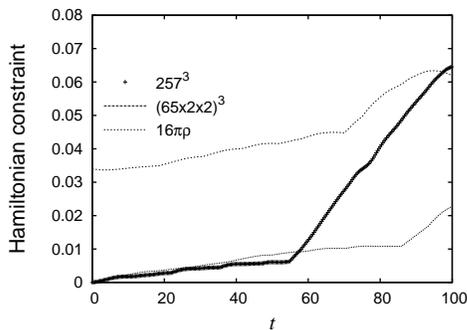
FIG. 7: HC violations for coalescing NSs in a $(65 \times 2 \times 2)^3$ AMR run compared to a 257^3 unigrid run. The maximum of $16\pi\rho$ (dotted line) is given for comparison.
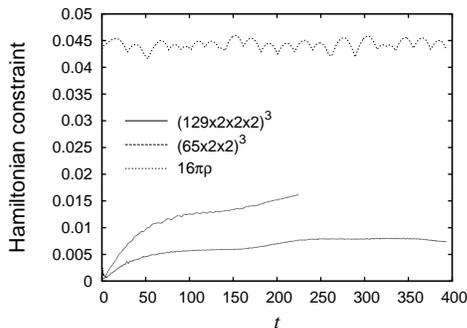
FIG. 8: Maximum HC violations for inspiraling NSs at two different resolutions, showing first order convergence. The maximum of $16\pi\rho$ (dotted line) is given for comparison.

veloped by the Meudon group [12, 13], and imported onto the Cartesian grid structure in GR-Astro-AMR for dynamical evolutions (again with Γ freezing shift and 1+log lapse).

We show results from two AMR simulations. In the high resolution run, the finest level grid has 60 points across the star, with 4 levels of refinement (refinement ratio is 2), and a coarse grid that is 129^3 covering a computational domain of $34R$. The low resolution run uses 3 levels of refinement and a coarse grid of 65^3, with the finest level having 30 points across each star. The maximum HC violations for the two runs are compared in fig. 8, where the maximum value of $16\pi\rho$ is also shown. We see that the error is converging to 1st order as expected. In fig. 9, the density on the equatorial plan in a gray scale plot with the grid structure imposed is shown at $t = 0$ and $t = 288M_\odot$ for the high resolution 4 level AMR run, with a downsampling factor of 4. We have zoomed in to show the inner part of the computational domain. At $t = 0$ the two finest grid patches are next to one another separated by a few grid points of the next level. The grid one further level down barely shows up with the down sampling. By $t = 288M_\odot$, the two NS has evolved for about 135 degrees, and the fine grid patches have merged to form something more complicated (in the

shape of 3 overlapping rectangles), while the two intermediate level grids becomes roughly squares. To provide another perspective of the grid structures, we show in fig. 10 the lapse as a height field at the corresponding times, with a downsampling factor of 2.

This study clearly shows the power of AMR. The higher resolution AMR simulation was carried out on a Dell desktop computer using 8 GB of memory. A regular unigrid simulation resolving the star and covering the computational domain to the same level as the AMR simulation would have to have 1025^3 grid points and require over 1.2TB of memory.

e. Discussions and Conclusions. We discussed the necessity of having the AMR capability in general relativistic simulations of NS coalescences. We demonstrated that the GR-Astro-AMR Code is capable of such AMR simulations. The simulations presented in this paper are the first steps towards what is needed in NS inspiral coalescence studies.

In sec. b of this paper we showed some of the validation studies we carried out for GR-Astro-AMR with a boosted NS. The simulation invokes all terms in the coupled Einstein-general relativistic hydro equations. Three different kinds of convergence tests have been performed, showing that the code has the convergence properties designed.

In sec. c we demonstrated that GR-Astro can be used to carry out NS coalescence simulations, and the accuracy is comparable to that of an unigrid simulation using a resolution same as that of the finest grid in the AMR run.

Sec. d showed an AMR simulation of an inspiraling NS binary carried out on a workstation. The simulation is equivalent in accuracy to a 1025^3 regular unigrid run, to the best of our knowledge larger than all previous simulations of similar systems on supercomputers. We also demonstrated that generation and merging of grid patches could be handled effectively in our AMR treatment.

In a future publication, we will extend the study in Sec. c to analyze the amount of matter available for accretion after the NS coalescence/BH formation, as a function of the angular momentum of the system at the plunge point of the inspiral. We will extend the study in Sec. d to determine astrophysically realistic initial data for inspiral, following the line initiated in [1]. These investigations require more computational resources than are available to us if they are to be carried out in unigrid.

There are many aspects in the GR-Astro-AMR code that need improvement as a computational infrastructure for a general relativistic simulation. In the next steps, we will (i) develop the parallel capacity of GR-Astro-AMR, (ii) study the usage of different refinement criteria for other NS/BH processes, (iii) develop a high order flux-conserving interpolation scheme for our hydro solver which uses a flux conservative treatment, and (iv) enable the direct solving of elliptic equations on the grid hierarchy.

(a) $t = 0$
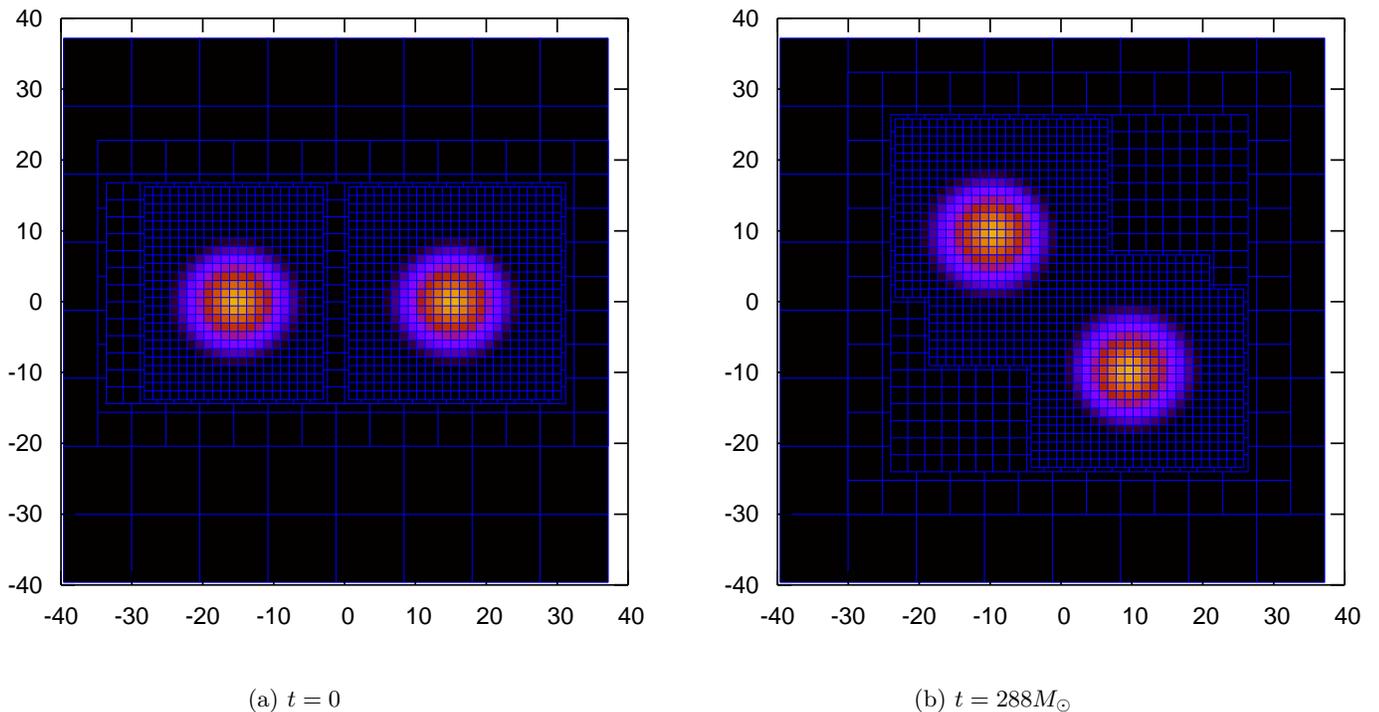
(b) $t = 288 M_\odot$

FIG. 9: Density for inspiralling neutron stars is given in grayscale with 4 levels of grid structure (downsampled by a factor of 4) superimposed. Only the central part of the $34R$ computational domain is shown. (a) $t = 0$, (b) $t = 288 M_\odot$ with the NSs rotated about $135°$ in their inspiral orbit.

(a) $t = 0$
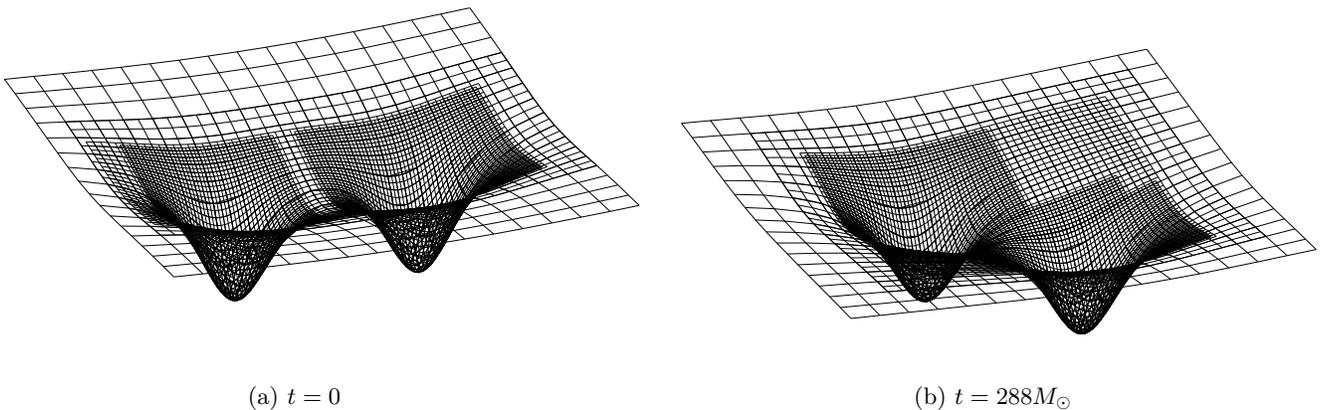
(b) $t = 288 M_\odot$

FIG. 10: The lapse with the grid structure superimposed in a different perspective (downsampled by 2). (a) $t = 0$, (b) $t = 288 M_\odot$. Only the central part is shown.

The code is developed with the intention of providing a computational tool to the general relativistic astrophysics community. The unigrid version of GR-Astro has been released (available at http://www.wugrav.wustl.edu). GR-Astro-AMR will be released as soon as ready. We invite researchers to join us in making use of as well as further developing this code.

[1] M. Miller and P. Gressman and W.-M. Suen, Phys. Rev. D **69**, 064026 (2004)

[2] Breno Imbiriba, John Baker, Dae-Il Choi, Joan Centrella, David R. Fiske, J. David Brown, James R. van Meter, Kevin Olson, Evolving a puncture black hole with fixed mesh refinement, gr-qc/0403048

[3] Erik Schnetter, Scott H. Hawley, Ian Hawke, Evolutions in 3D numerical relativity using fixed mesh refinement, Class.Quant.Grav. 21 (2004) 1465-1488, gr-qc/0310042

[4] Dae-Il Choi, J. David Brown, Breno Imbiriba, Joan Centrella, Peter MacNeice, Interface Conditions for Wave Propagation Through Mesh Refinement Boundaries, J.Comput.Phys. 193 (2004) 398-425. physics/0307036

[5] Frans Pretorius, Luis Lehner, Adaptive Mesh Refinement for Characteristic Codes, J.Comput.Phys. 198 (2004) 10-34. gr-qc/0302003

[6] Kimberly C. B. New, Dae-Il Choi, Joan M. Centrella, Peter MacNeice, Mijan F. Huq, Kevin Olson, Three-dimensional adaptive evolution of gravitational waves in numerical relativity, Phys.Rev. D62 (2000) 084039, gr-qc/0007045

[7] Peter Diener, Nina Jansen, Alexei Khokhlov, Igor Novikov, Adaptive mesh refinement approach to construction of initial data for black hole collisions, Class.Quant.Grav. 17 (2000) 435-451, gr-qc/9905079

[8] http://www.caip.rutgers.edu/TASSL/Projects/GrACE/

[9] http://wugrav.wustl.edu/research/projects/nsgc.html

[10] http://www.cactuscode.org

[11] J. A. Font, Numerical hydrodynamics in general relativity, Living Reviews in Relativity, 6, (2003), 4, http://www.livingreviews.org/lrr-2003-4

[12] E. Goulgoulhon et al., Quasiequilibrium sequences of synchronized and irrotational binary neutron stars in general relativity: Method and tests, Phys. Rev. D63, 064029, (2001)

[13] K. Taniguchi and E. Gourgoulhon, Quasiequilibrium sequences of synchronized and irrotational binary neutron stars in general relativity. III. Identical and different mass stars with gamma=2, Phys. Rev. D66, 104019, (2002)

[14] Mari Kawamura, Ken-ichi Oohara, Takashi Nakamura, General Relativistic Numerical Simulation on Coalescing Binary Neutron Stars and Gauge-Invariant Gravitational Wave Extraction, astro-ph/0306481.

[15] Masaru Shibata, Keisuke Taniguchi, Koji Uryu, Merger of binary neutron stars of unequal mass in full general relativity, Phys.Rev. D68 (2003) 084020, gr-qc/0310030.

[16] Joshua A Faber, Philippe Grandclment, Frederic A Rasio, Mergers of Irrotational Neutron Star Binaries in Conformally Flat Gravity, gr-qc/0312097.