

# Astrophysical Applications of Numerical Relativity — from Teragrid to Petascale

Erik Schnetter<sup>(1,2)</sup>, Christian D. Ott<sup>(3)</sup>, Peter Diener<sup>(1,2)</sup>, Christian Reisswig<sup>(4)</sup>

<sup>(1)</sup> Center for Computation & Technology, Louisiana State University, Baton Rouge, LA, USA, <http://www.cct.lsu.edu/>

<sup>(2)</sup> Department of Physics and Astronomy, Louisiana State University, Baton Rouge, LA, USA, <http://www.phys.lsu.edu/>

<sup>(3)</sup> Department of Astronomy and Steward Observatory, The University of Arizona, Tucson, AZ, USA, <http://www.as.arizona.edu/>

<sup>(4)</sup> Max-Planck-Institut für Gravitationsphysik (Albert-Einstein-Institut), Potsdam, Germany, <http://www.aei.mpg.de/>

## ABSTRACT

The TeraGrid is growing to a Petascale environment. Its latest additions and future planned systems will allow astrophysicists to investigate astrophysical objects that cannot be studied with today's observational methods, and which were previously excluded from computational study by lack of both computational power and appropriate codes. We present an overview of recent results in relativistic astrophysics, including binary black hole systems and stellar core collapse scenarios. We describe the main properties of our software infrastructure and our physics codes that we employ on the TeraGrid, and we outline our plans towards new and groundbreaking physics in the coming years.

## 1 INTRODUCTION

Since the invention of the microprocessor, there has been an exponential increase in the (theoretical peak) single CPU performance and consequently also in the integral peak performance of highly parallel supercomputers. Machines reaching the petaflop mark are being designed and built; they will be more than four million times faster than the Cray-1 which still serves as the icon of supercomputing.

However, computer software is today not able to keep up with harnessing this wealth of computing power. In particular, many physics codes written for the Cray-1 thirty years ago are still being used virtually unaltered on today's workstations or as parts of large simulation codes running on modern high performance systems. These codes are typically not well optimized for modern architectures, leading to severe limitations on efficient resource usage such as the observed large gap between theoretical peak and actually achieved floating point performance or parallel scaling.

Although hundreds of millions of dollars are being targeted at designing and deploying petascale hardware and at developing new programming languages and models, there is still a lack of understanding of how they will be employed for real-life and real-science applications [8].

One such science problem which is of interest to us is the merging of two black holes. This is expected to be one of the strongest sources of gravitational radiation in the universe. Gravitational waves have not been observed directly to date, but gravitational wave detectors (e.g. the NSF-funded LIGO [1]) are in the process of reaching sensitivities

sufficiently high to observe interesting astrophysical phenomena. Another, particular challenging problem in astrophysics – the modeling of gamma-ray bursts (GRBs, introduced in section 4.2 below) – is motivating us to be ready to exploit petascale computers, and is driving our physics and computer science development program.

Although solving this GRB simulation problem is likely to be a decade-long challenge, we believe that continuing to develop our scientific software through the component framework Cactus [2, 37] will provide a path to petascale. In this paper we present an overview of recent results in two particular areas of relativistic astrophysics, namely binary black hole systems and stellar core collapse, which are stepping stones towards for GRB simulations. We further describe the main properties of our software infrastructure and our physics codes, and then outline our plans for the coming years.

## 2 METHOD DESCRIPTION

### 2.1 PHYSICS

We carry out our astrophysical simulations in full non-linear general relativity (GR) without any symmetry assumptions (i.e., in 3D) using the 3+1 formalism of Einstein's GR (see e.g. [17] for a review) that foliates spacetime into time-consecutive spacelike 3-hypersurfaces. Einstein's equations are treated as a Cauchy initial value problem and are evolved freely. In the particular formalism we employ – BSSN (see e.g. [9, 10, 16, 17, 20] for details on the formulation and numerical implementation) – the Einstein equations separate into 17 evolution equations, 8 gauge evolution equations, and 4 constraint equations. The latter must be fulfilled by the initial data and continued fulfillment is guaranteed by the evolution equations in the continuum limit. In a numerical evolution at finite resolution the magnitude of the constraint violations is used (1) to evaluate the quality of a simulation and (2) for testing proper convergence of the code. We employ our legacy curvature evolution code CCATIE [3, 9, 10] and the newly developed state-of-the-art code McLachlan for the astrophysically highly relevant scenario of binary black-hole inspiral and merger (see e.g. [41, 55, 56, 57, 58]).

For relativistic astrophysical systems that contain dynamically-relevant amounts of ordinary matter, GR hydrodynamics is a crucial second ingredient, providing, in coupled fashion with the curvature part, for the time-update of the matter fields (for a review and details on GR hydrodynamics see [34]). For this we employ the publicly available finite-volume high-resolution shock capturing code *Whisky* [11]. We are also in the process of developing the GR magneto-hydrodynamics (GRMHD) code *DiFranco* [62] which will allow us to include the effects of magnetic fields in our simulations. Most of our GR hydrodynamics efforts have been focused on studying the collapse of massive stars to neutron stars [50, 52, 53], the collapse of neutron stars to black holes [11, 12, 14], non-axisymmetric rotational instabilities in neutron stars [13, 52], neutron star binaries [30], as well as neutron-star – black hole binary systems [44].

## 2.2 COMPUTATIONAL INFRASTRUCTURE

We structure our code through the Cactus framework [2, 37], which lets us compose individually developed components (called *thorns*) to full applications. Cactus is a tightly coupled framework, intended to run within one high-bandwidth low-latency system. Components share their main data structures, while memory, communication, and I/O are managed by a specialised *driver* component. This architecture enables highly efficient component coupling with virtually no overhead, but imposes certain restrictions onto the architecture of the individual components.

The framework itself (called *flesh*) does not offer any significant functionality on its own. It rather offers a set of well-designed APIs which are implemented by other components; the specific set of components which are used to implement these can be decided at run time. Such APIs provide coordinate systems, generic interpolation, reduction operations, hyperslabbing, and various I/O methods, and more.

Cactus originated in the numerical relativity community in 1997, and has since matured into a generic software framework for high-end computing applications which are developed by large and international collaborations. Cactus is used by more than two dozen research groups worldwide and is part of the petascale tool development for NSF’s Blue Waters Petascale Computing System [4].

The specific physics equations and their discretization are implemented as thorns. We discretize the simulation domain using high-order finite differences on block-structured grids, employing Berger-Oliger adaptive mesh refinement (AMR) [19]. We use explicit Runge-Kutta type time integration methods, and subcycling in time on finer grids ensures a constant CFL factor independent of the spatial resolution. Our AMR algorithm is implemented in the *Carpet* driver component [5, 60] which provides the basic recursive AMR algorithm. Carpet builds on a lower level *CarpetLib* component providing memory management and inter-process communication, and is accompanied by helper components providing I/O, grid structure setup, and other functionality.

In addition to mesh refinement, Carpet supports multi-patch systems [29, 59, 62] where the domain is covered by multiple, possibly overlapping, distorted, but logically Cartesian grid blocks.

The framework architecture, where each component declares its interfaces and where memory is managed by a specialised driver component, is akin to function composition in functional programming languages. Components do not store any state themselves, they only act on variables which are passed in by the framework. This has a number of advantages: debugging is simplified since the data flow is known explicitly, all variables can be easily output since the driver manages them, and checkpointing/recovery is almost trivial to implement since the driver knows the complete state of the simulation at all times. This assumes that components play by the rules – this is not actually enforced, but components have to manage exceptions on their own.

Cactus parallelizes its data structures on distributed memory architectures via spatial domain decomposition, with *ghost zones* added to each MPI process’ part of the grids. Synchronisation is performed implicitly, by declaring for each routine which variables it modifies, instead of by explicit calls to communication routines.

Higher order methods require a substantial amount of ghost zones (in our case three ghost zones for fourth order accurate, upwinding differencing stencils), leading to a significant memory overhead for each MPI process. This can be counteracted by using OpenMP [6] within a multi-core node, which is especially attractive on modern systems with eight (Abe, Queen Bee) or 16 (Ranger) cores per node; all performance critical parts of Cactus support OpenMP. However, NUMA systems such as Ranger require care in laying out data structures in memory to achieve good OpenMP performance, and we are therefore using a 4-by-4 combination on this system (4 MPI processes with 4 OpenMP

threads each).

We have recently used *Kranc* [7, 40, 42] to generate a new BSSN-type Einstein code *McLachlan* from scratch. Kranc is a Mathematica-based code generation system that starts from continuum equations entered into Mathematica, and automatically generates full Cactus thorns after discretizing the equations. This approach shows a large potential, not only for reducing errors in complex systems of equations, but also for reducing the time to implement new discretization methods such as higher-order finite differencing or curvilinear coordinate systems. It furthermore enables automatic code optimisations at a very high level, using domain-specific knowledge about the system of equations and the discretization method that is not available to the compiler, such as cache or memory hierarchy optimisations or multi-core parallelization. Such optimisations are planned for the future.

## 3 RESOURCE USAGE

We regularly use the TeraGrid resources Abe, Queen Bee, or Ranger in addition to local resources for large-scale science runs. We are happy to mention that we managed to increase the efficiency and scalability of our computational infrastructure during the friendly-user periods of each of these machines; we find that the noncommittal way in which large amounts of CPU time can be used (or even abused) for tests and the short queue turn-around times form an ideal playground on which new ideas can be implemented and tried with relative ease.

In a “typical” medium-size production run we use, in addition to a series of smaller companion runs, about 32–64 CPU cores with about 2 GByte of memory per core for binary black hole (BBH) simulations; such runs may last for several days of wall clock time, for several tens of kSU per simulation. A large-size production run uses significantly more resources. Until recently, we were unfortunately restricted to using not more than about 256 cores per run due to inefficiencies in our communication infrastructure, but several “hero runs” used 64 cores for about 90 days of wall time [50, 52] to solve a particularly demanding stellar-collapse problem.

In more recent production runs for our binary black hole simulations, we use between 128 and 256 cores, since some observables of astrophysical interest require higher resolutions and larger simulation domains to achieve the accuracy necessary for our results. As parameter studies, i.e., the study of many different BBH models, are indispensable in our work [56, 57, 58], we easily reach the capacities of terascale clusters.

Figure 1 shows results of weak scaling tests of the McLachlan code on Franklin (NERSC), Queen Bee (LONI), and Ranger (TACC), several large machines that have recently become available. The machines’ configuration is briefly described in table 1. Shown is the number of grid point updates that are calculated per millisecond where larger is better. The computational load per core was kept constant, increasing the total problem size with the number of cores. Ideal scaling would be a horizontal line. For large numbers of cores, the AMR infrastructure’s efficiency drops especially on Franklin, but simulations with up to about 8192 cores on Ranger have reasonable efficiency. The exact reason for the non-scaling is currently under investigation.

## 4 PHYSICS DRIVERS AND RESULTS

### 4.1 BINARY BLACK HOLE MERGERS

Binary black hole (BBH) mergers are expected to be the strongest sources of gravitational radiation in the universe and are therefore of main interest to the currently active first generation of gravitational wave detectors. Hence, it is an important task for Numerical Relativity to produce gravitational wave templates for various different

Resource	Host	Type	Arch.	CPU Clock	Theor. Peak	Interconnect	Nodes	Cores	C/N
Franklin	NERSC	Cray XT4	AMD	2.6 GHz	5.2 Gflop/s	Hypertransport	9432	18864	2
Ranger	TACC	Sun Constellation	AMD	2.0 GHz	8.0 Gflop/s	Infiniband	3936	62976	16
Queen Bee	LONI	Dell	Intel	2.33 GHz	9.32 Gflop/s	Infiniband	668	5344	8

Table 1: Overview of the machines used for the benchmarks shown in figure 1. C/N denotes the number of cores per node.

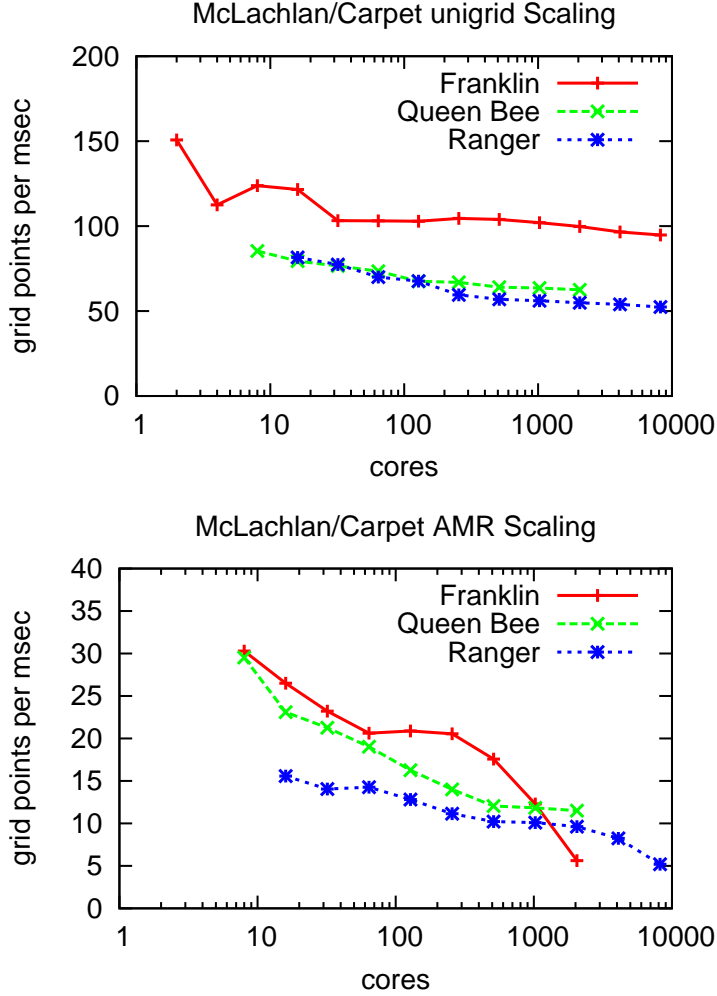


Figure 1: Weak scaling results of the McLachlan code for unigrid and for mesh refinement setups, evolving the full Einstein equations on large contemporary systems. Shown is the number of grid point updates calculated per millisecond (larger is better). These tests employed a hybrid MPI/OpenMP communication model (except on Franklin). Even with mesh refinement, there is reasonable scaling to up to more than 8000 cores.

BBH models to sufficiently cover the parameter space of possible detections. These templates can be used within gravitational data analysis in matched filtering searches to increase the probability in detecting a wave signal for such events.

The initial data for our simulations are chosen to be on approximately quasi-circular orbits.<sup>1</sup> We typically use 8 levels of mesh refinement in order to have high resolution only where it is needed, i.e., close to the black holes where the curvature is largest. Figure 2 shows the initial configuration and grid structure for one of our simulations. Only the 5 finest levels of refinement are visible in the figure; the others ex-




Figure 2: Initial configuration for a binary black hole merger, showing both the location and size of the black holes (black) and the grid structure of our mesh refinement setup in the  $z = 0$  plane. The height field is the lapse function, one of the metric components indicating the curvature of the spacetime. Curvature becomes infinite near the centers of the black holes.

tend further outwards. As the black holes move across the domain, the refined regions move to track them.

The tracks of the black holes in a typical simulation are shown in figure 3. As the black holes emit gravitational waves, the orbit shrinks and the black holes speed up. This in turn leads to a higher rate of gravitational wave emission and even faster shrinking of the orbit. The end result is the spiral pattern of figure 4, which shows a measure of the gravitational wave content on an equatorial slice shortly after the merger. Since gravitational waves propagate with the speed of light, a larger radial coordinate corresponds to an earlier emission time. This leads to a chirp-like gravitational wave signal as shown in figure 5 where both the amplitude and frequency increases with time until the two black holes merge into one. This resulting black hole is initially highly distorted, and it settles down to a stationary state through highly damped oscillations. In this phase the gravitational wave signal is essentially an exponentially damped sinusoid.

Several important questions of astrophysical relevance can be addressed by BBH simulations. Among those is the recoil velocity that the merger remnant acquires due to the asymmetric emission of gravitational radiation. Gravitational waves carry away a total net linear momentum during the inspiral and merger, and due to the conservation of linear momentum the merger remnant ends up obtaining a gravitational recoil. This recoil or “kick” [18, 54] can be of astrophysical significance in different ways. In the case of supermassive black holes, this may lead to an understanding of how our universe evolved during its early stages of galaxy formation in the epoch of structure formation through hierarchical galaxy mergers. It is likely that each galaxy hosts a central supermassive black hole, and these will merge during galaxy mergers. If the kick of the merged central black hole is sufficiently large, it may overcome the binding energy of the merged

<sup>1</sup>If no gravitational waves were emitted the black holes would move on circular orbits.




Figure 3: Tracks of two equal-mass black holes with unequal spin as they spiral in towards each other. The tracks end when the black holes merge to a single black hole. It is clearly visible that the “inspiral speed” increases as the black holes approach each other. Figure 5 shows the gravitational waveform generated by such an event.

host galaxy so that it is ejected. The host would subsequently follow a different evolutionary path compared to galaxies with a central black hole [43, 45, 46].

In the case of stellar mass black holes, the same consideration applies to the merger of globular clusters within our own galaxy and may give us a hint on the distribution of these black holes in our galactic disk [45, 48, 49].

While estimates of kick velocities have been available for some time [31, 32, 33], the largest part of the system’s acceleration is generated in the final orbit and plunge of the binary system, and, thus, requires fully-relativistic calculations to be determined accurately, and much progress was recently made in this direction [15, 23, 27, 28, 35, 36, 38, 39, 41].

We have performed a series of simulations computing kick velocities for different initial configurations. The parameter space of initial configurations is seven-dimensional; there are three components for the initial spin for each of the black holes as well as the mass ratio. We restricted ourselves to equal-mass black holes whose spin is aligned with the orbital angular momentum leaving us with only two free parameters: the  $z$ -component of the spins  $a_1, a_2$  for each of the two black holes. This is motivated by evidence for a preference in nature for binaries that have spins aligned with the orbital angular momentum [21], at least in gas-rich mergers.

In our simulations we found a quadratic dependence of the kick velocity on the initial spins  $a_1, a_2$ . Furthermore, in our simulations the kick is not sufficient to lead to an ejection from the host object [41, 55, 58]. This may be an artifact of the restricted parameter space which we examined. Other works with different initial spin configurations led to so called “super-kicks” [23, 27, 28] which may easily exceed the escape velocity of average galaxies. However, these configurations are less likely to occur in nature [21].

A further astrophysically interesting result is the dependence of the spin of the final black hole on the initial spins of the individual black holes. We were able to obtain a fitting formula for the final spin as a function of the initial spin parameters, not only for spin-aligned binaries [57, 58] with equal mass, but also for arbitrarily spinning black holes with unequal masses [56]. Both of these phenomenological formulas for the kick velocity and for the spin magnitude can be readily used in  $N$ -body simulations of stellar cluster and galaxy simulations.




Figure 4: Gravitational radiation emitted in the orbital plane during a BBH merger encoded in the Weyl scalar  $\Psi_4$ . Plotted here is  $r \cdot \Psi_4$  to compensate for the  $1/r$  falloff that  $\Psi_4$  obeys. One can see that the amount of radiation is particularly strong during a certain time - the time of the merger - and becomes weaker in the ring-down phase where the final black hole settles to a stationary state. The remnant can not be seen at this scale since its horizon radius is around  $r_H \approx 0.3M$ . The emitted radiation is slightly asymmetric.

In these cases, the numerical solution of the full Einstein equations in black hole encounters is unfeasible and has to be included by approximate formulas.

## 4.2 STELLAR CORE COLLAPSE

Stars more massive than  $\sim 10$  solar masses ( $M_\odot$ ) burn their nuclear fuel all the way to the peak of the nuclear binding energy curve marked by nuclei of the iron group. The central iron core is surrounded by an onion-skin structure composed of ashes of previous burning stages. The iron core eventually loses its central pressure support and collapses to a protoneutron star (PNS). Note that collapse never proceeds directly to a black hole, since nuclear repulsive forces always lead to a rebound-like (“core bounce”) stabilization at the protoneutron star stage. Black hole formation can only occur through mass accretion at later times.

The collapse to a protoneutron star releases a prodigious amount of gravitational energy, a fraction of which could power a core-collapse supernova, the dramatic event marking massive star death (for a semi-popular review, see [25]).

The fundamental problem that supernova theory has faced in the past 50 years is how exactly the necessary fraction (about a hundredth) of gravitational energy is converted into the energy of the supernova explosion. Although significant progress has been made, the community is far from agreeing on the supernova mechanism [61] and the core-collapse supernova problem remains unsolved to the present day.

A second, perhaps even more pressing, unsolved problem associated with core-collapse supernovae is their detailed relationship with gamma-ray bursts (GRBs; see e.g. [47] for a review). GRBs are intense narrowly-beamed flashes of  $\gamma$ -rays (very high-energy photons) of cosmological origin. They come in two groups: long-soft bursts last between  $\sim 10$ –1000 seconds and have  $\gamma$ -spectra peaking at lower fre-






Figure 5: Gravitational waveform generated by a black hole merger event. As the black holes orbit each other, both frequency and amplitude of the wave increases, peaking as a single black hole is formed. The single black hole is initially highly distorted and rings down quickly. Shown is the “plus”-polarization of the gravitational wave strain.

quencies than those of short-hard bursts, that have durations of  $\sim 2$  seconds. Long-soft GRBs show features in their longer-wavelength afterglow emissions that strongly suggest a close relationship with very energetic type-Ic core-collapse supernovae from massive stars that have lost their hydrogen envelope and a fair fraction of their helium shell to a stellar wind ( $M \gtrsim 25\text{--}35 M_\odot$ ). Just as the supernova mechanism, the mechanism powering GRBs is currently not known. Energetics and phenomenology suggests that GRBs may be powered by a central engine, consisting of a stellar-mass black hole that accretes from a rotationally-supported accretion disk. The rotational and gravitational energy of the accreting material could be converted by this engine (via magneto-rotational effects and/or by the annihilation of neutrino – anti-neutrino pairs) into the kinetic energy of a bipolar relativistic outflow (jet) that punches through the stellar envelope and produces the characteristic  $\gamma$ -ray emission.

The high-density, strong-gravity regions in which the engines of both GRBs and core-collapse supernovae operate are impossible to observe by the means of conventional astronomy. Electromagnetic radiation, the basic messenger of astronomy, cannot propagate through the intervening stellar material. Only neutrinos and gravitational waves have the potential of acting as live messengers for the intricate dynamics associated with GRB and core-collapse supernova central engines. In order to ascertain central-engine details, numerical models are required for predicting and matching gravitational wave and neutrino observations.

Realistic numerical models of stellar collapse – the presumed precursor to both core-collapse supernovae and long-soft GRBs – must encompass a multitude of physics, including GR, GRMHD, neutrino transport and interactions, and a finite-temperature nuclear equation of state, relating density, temperature and composition to pressure and other thermodynamic quantities. In addition, simulations should be carried out without symmetry restrictions, hence in 3D, in order to capture all possible multi-dimensional dynamics/phenomena.

In our efforts towards simulating core-collapse supernovae and GRBs, we have recently reached an important milestone by performing the first self-consistent 3D simulations of massive star collapse in full non-linear GR [50, 52, 53]. These simulations also included for the first time all relevant microphysics necessary (neutrino treatment, finite-temperature nuclear equation of state) to model realistic young




Figure 6: Volume rendering of the rest-mass density of a rotationally-deformed protoneutron star (PNS) shortly after core bounce in a 3D GR hydrodynamics simulation carried out by Ott et al. [50, 52, 53]. The pattern displayed on the plane below the PNS reflects the Weyl scalar  $\Psi_4$ , representing qualitatively the outgoing gravitational wave field.




Figure 7: Gravitational wave signature of stellar collapse, bounce, and early post-bounce phase in a rapidly-spinning 3D GR model by Ott et al. [50, 52, 53]. Realistic cores stay axisymmetric through core bounce and the very early post-bounce phase. This is reflected in the absence of gravitational wave emission along the axis of symmetry. At later times, non-axisymmetric structures and dynamics develop, leading to prolonged and strong emission of gravitational waves also along the polar direction.

protoneutron stars. We performed this calculations for the 200–300 millisecond initial collapse phase and followed the protoneutron-star – outer iron core system to  $\sim 100$  ms after core bounce. Our focus in this study was on the dynamics of collapse and bounce and on the development of 3D structures and dynamics at post-bounce times.

The results of our simulations show that stellar core collapse proceeds in rotational symmetry (axisymmetry) through collapse and bounce. At post-bounce times, however, the protoneutron star may undergo a rotational hydrodynamic instability that drives the development of non-axisymmetric structure as depicted in figure 6 for a particularly rapidly and differentially spinning protoneutron star simulation. Figure 7 displays the gravitational waves emitted in a representa-

tive simulation of rapidly-rotating stellar core collapse. A strong burst of gravitational waves is emitted at bounce when the inner core is most compact and the accelerations are greatest. What follows are damped ring-down oscillations in which the newly-formed protoneutron star dissipates its remaining pulsational energy. Core bounce and ring down occur in axisymmetry and are in gravitational waves visible only to observers situated away from the axis of symmetry. At post-bounce times  $\gtrsim 20\text{--}30$  ms, the protoneutron star develops non-axisymmetric substructure in regions near its equatorial plane, leading to the emission of quasi-periodic gravitational waves (blue graph in figure 7; as seen by observers situated along the poles of the system). Although of lower amplitude, the prolonged duration of this post-bounce emission can easily dwarf the bounce signal in total energy emission and is an important enhancement of the stellar collapse gravitational wave signature.

## 5 FUTURE IMPROVEMENTS

We assume that our present spacetime and GR (magneto-)hydrodynamics codes can handle self-consistently and with good physical detail both the initial collapse and early post-bounce phase of core-collapse supernovae [50, 52, 53] and the dynamics of black hole formation and subsequent accretion [51].

To model with sufficient physical detail the workings of core-collapse supernova and GRB central engines, the inclusion of neutrino transport and interactions is crucial. This will ultimately require radiation transport in a multi-group flux-limited diffusion approximation. Building upon experience with transport solvers in the VULCAN/2D code (see, e.g., [26]), this could be implemented in an efficient ray-by-ray fashion (e.g. [24]). Solving a diffusion equation at every time step as well as the necessary implicit time integration schemes will necessitate efficient solvers, increasing the computational requirements of our code by at least an order of magnitude.

These plans for additional physics need to be accompanied by corresponding improvements in our computational infrastructure. Also, future architectures such as Blue Waters will require unprecedented levels of parallelism from codes that intend to achieve petascale performance.

We expect some benefit from additional static optimizations, implemented, e.g., within the Kranc code generator, including generating explicit vector instructions for inner loops, eliminating common subexpressions, and reducing the amount of temporary storage. In our experience, compilers do not perform well at these tasks for our large inner loops. Automatic code generation will let us apply these by re-generating affected code. Another set of optimizations can be applied dynamically, i.e., at run time depending on the particular system hardware; examples are loop tiling, prefetching, and data copying, or changes to the memory layout of our 3D grid functions to optimize cache efficiency. Using a framework allows us to perform experiments in this direction without rewriting the user code implementing the equations or containing the discretization. All such optimizations should be automatically evaluated at run time, both to give the user feedback on their efficiency as well as to collect a performance database for future research.

The GRB simulations will generate large amounts of output data, both in size and in the number of simulations that will be run. Metadata are central to the management, archival and retrieval of simulation data. In the metadata model currently under development for Cactus, we define for each simulation extensible key-value pairs of *core* and *key* metadata. *Core* metadata describe the basic and generic characteristics of the simulation, including parameter settings, information on active Cactus thorns, size of the simulation, machine and performance information, source code version etc. *Key* metadata, on the other hand, consist of application-specific key-value pairs, for ex-

ample, for a GRB simulation, they may contain a description of the progenitor star characteristics (mass/rotational configuration etc.), information on when the black hole horizon first appears, or when the GRB jet breaks through the stellar envelope. Each core and key metadata key-value pair may be classified static (cannot change during the simulation) or dynamic (can change during the simulation). In this way, the simulation metadata can also be used for on-line remote monitoring of central characteristics while a simulation is running [22].

With the additional physics capabilities described above, and with an improved and extended computational infrastructure, we hope to be able to use current and future TeraGrid resources efficiently to make good progress towards new and groundbreaking physics results.

## ACKNOWLEDGEMENTS

The results presented in this paper use the Cactus framework, the Carpet adaptive mesh refinement infrastructure, the CCATIE spacetime code, and the Whisky general relativistic hydrodynamics package. We thank their respective authors and maintainers. We also thank G. Allen, F. Löffler, and J. Tao for help with defining Cactus benchmarks.

C.D.O. acknowledges support through a Joint Institute for Nuclear Astrophysics postdoctoral fellowship, sub-award no. 61-5292UA of NFS award no. 86-6004791. C.R. acknowledges support through the International Max Planck Research School fellowship. This research employed the resources of the Center for Computation & Technology at Louisiana State University, which is supported by funding from the Louisiana legislature’s Information Technology Initiative. We acknowledge support for the XiRel project via the NSF PIF award 0701566 and for the Alpaca project via the NSF SDCI award 0721915. This research was supported in part by resources provided by the Louisiana Optical Network Initiative, by the National Science Foundation through TeraGrid resources provided by LONI, NCSA, and TACC, by the National Energy Research Scientific Computing Center, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231, and by the resources of the Max-Planck-Institute for Gravitational Physics (AEI).

## REFERENCES

- [1] LIGO: Laser Interferometer Gravitational Wave Observatory, URL <http://www.ligo.caltech.edu/>.
- [2] Cactus Computational Toolkit home page, URL <http://www.cactuscode.org/>.
- [3] Spacetime evolution with CCATIE, URL <http://numrel.aei.mpg.de/Research/codes.html>.
- [4] Blue Waters, URL <http://www.ncsa.uiuc.edu/BlueWaters/>.
- [5] Mesh refinement with Carpet, URL <http://www.carpetcode.org/>.
- [6] OpenMP: Simple, Portable, Scalable SMP Programming, URL <http://www.openmp.org/>.
- [7] Kranc: Automated Code Generation, URL <http://numrel.aei.mpg.de/Research/Kranc/>.
- [8] HPC Grand Challenge Case Study, Keeping the Lifeblood Flowing: Boosting Oil and Gas Recovery from the Earth, March 2005. Sponsored by DOE, URL <http://www.compete.org/publications/detail/389/grand-challenge-case-study-oil-and-gas-recovery/>.

- [9] M. Alcubierre, B. Brügmann, P. Diener, M. Koppitz, D. Pollney, E. Seidel, and R. Takahashi. Gauge conditions for long-term numerical black hole evolutions without excision. *Phys. Rev. D*, 67:084023, 2003. eprint gr-qc/0206072.
- [10] M. Alcubierre, B. Brügmann, T. Dramlitsch, J. A. Font, P. Papadopoulos, E. Seidel, N. Stergioulas, and R. Takahashi. Towards a stable numerical evolution of strongly gravitating systems in general relativity: The conformal treatments. *Phys. Rev. D*, 62:044034, 2000. eprint gr-qc/0003071.
- [11] L. Baiotti, I. Hawke, P. J. Montero, F. Löffler, L. Rezzolla, N. Stergioulas, J. A. Font, and E. Seidel. Three-dimensional relativistic simulations of rotating neutron star collapse to a Kerr black hole. *Phys. Rev. D*, 71:024035, 2005. eprint gr-qc/0403029.
- [12] L. Baiotti, I. Hawke, L. Rezzolla, and E. Schnetter. Gravitational-wave emission from rotating gravitational collapse in three dimensions. *Phys. Rev. Lett.*, 94:131101, 2005. eprint gr-qc/0503016.
- [13] L. Baiotti, R. D. Pietri, G. M. Manca, and L. Rezzolla. Accurate simulations of the dynamical bar-mode instability in full general relativity. *Phys. Rev. D*, 75:044023, 2007. eprint [http://arXiv.org/abs]astro-ph/0609473.
- [14] L. Baiotti and L. Rezzolla. Challenging the paradigm of singularity excision in gravitational collapse. *Phys. Rev. Lett.*, 97:141101, 2006. eprint gr-qc/0608113.
- [15] J. G. Baker, J. Centrella, D.-I. Choi, M. Koppitz, J. van Meter, and M. C. Miller. Getting a kick out of numerical relativity. *Astrophys. J.*, 653:L93–L96, 2006. eprint astro-ph/0603204.
- [16] T. W. Baumgarte and S. L. Shapiro. On the numerical integration of Einstein’s field equations. *Phys. Rev. D*, 59:024007, 1999. eprint gr-qc/9810065.
- [17] T. W. Baumgarte and S. L. Shapiro. Numerical relativity and compact binaries. *Physics Reports*, 376(2):41–131, March 2003. eprint gr-qc/0211028.
- [18] J. D. Bekenstein. Extraction of energy and charge from a black hole. *Phys. Rev.*, D7:949–953, 1973.
- [19] M. J. Berger and J. Olinger. Adaptive mesh refinement for hyperbolic partial differential equations. *J. Comput. Phys.*, 53:484–512, 1984.
- [20] H. Beyer and O. Sarbach. On the well posedness of the Baumgarte-Shapiro-Shibata-Nakamura formulation of Einstein’s field equations. *Phys. Rev. D*, 70:104004, 2004. eprint gr-qc/0406003.
- [21] T. Bogdanovic, C. S. Reynolds, and M. C. Miller. Alignment of the spins of supermassive black holes prior to coalescence. 2007. eprint astro-ph/0703054.
- [22] R. Bondarescu, G. Allen, G. Daues, I. Kelley, M. Russell, E. Seidel, J. Shalf, and M. Tobias. The astrophysics simulation collaboratory portal: a framework for effective distributed research. *Future Generation Computer Systems*, 2003. Accepted.
- [23] B. Brügmann, J. A. González, M. Hannam, S. Husa, and U. Sperhake. Exploring black hole superkicks. 2007. arXiv:0707.0135 [gr-qc].
- [24] R. Buras, M. Rampp, H.-T. Janka, and K. Kifonidis. Two-dimensional hydrodynamic core-collapse supernova simulations with spectral neutrino transport. I. Numerical method and results for a 15  $M_{\odot}$  star. *Astron. Astrophys.*, 447:1049, 2006.
- [25] A. Burrows. Supernova Explosions in the Universe. *Nature*, 403:727, 2000.
- [26] A. Burrows, E. Livne, L. Dessart, C. D. Ott, and J. Murphy. Features of the Acoustic Mechanism of Core-Collapse Supernova Explosions. *Astrophys. J.*, 655:416, 2007.
- [27] M. Campanelli, C. O. Lousto, Y. Zlochower, and D. Merritt. Large merger recoils and spin flips from generic black-hole binaries. *Astrophys. J.*, (659):L5–L8, 2007. eprint gr-qc/0701164.
- [28] M. Campanelli, C. O. Lousto, Y. Zlochower, and D. Merritt. Maximum gravitational recoil. *Phys. Rev. Lett.*, (98):231102, 2007. eprint gr-qc/0702133.
- [29] P. Diener, E. N. Dorband, E. Schnetter, and M. Tiglio. Optimized high-order derivative and dissipation operators satisfying summation by parts, and applications in three-dimensional multi-block evolutions. *J. Sci. Comput.*, 32:109–145, 2007. eprint gr-qc/0512001.
- [30] E. Evans, S. Iyer, E. Schnetter, W.-M. Suen, J. Tao, R. Wolfmeyer, and H.-M. Zhang. Computational relativistic astrophysics with adaptive mesh refinement: Testbeds. *Phys. Rev. D*, 71:081301(R), 2005. eprint gr-qc/0501066.
- [31] M. Favata, S. A. Hughes, and D. E. Holz. How black holes get their kicks: Gravitational radiation recoil revisited. *Astrophys. J.*, 607:L5–L8, 2004. eprint astro-ph/0402056.
- [32] M. J. Fitchett. *Mon. Not. R. astr. Soc.*, 203:1049, 1983.
- [33] M. J. Fitchett and S. Detweiler. Linear momentum and gravitational waves – Circular orbits around a Schwarzschild black hole. *Mon. Not. R. Astron. Soc.*, 211:933–942, Dec. 1984.
- [34] J. A. Font. Numerical hydrodynamics in general relativity. *Living Rev. Relativity*, 6:4, 2003. URL <http://www.livingreviews.org/Articles/lrr-2003-4>.
- [35] J. A. Gonzalez, M. D. Hannam, U. Sperhake, B. Brügmann, and S. Husa. Supermassive kicks for spinning black holes. *Phys. Rev. Lett.*, 98:231101, 2007. eprint gr-qc/0702052.
- [36] J. A. González, U. Sperhake, B. Brügmann, M. Hannam, and S. Husa. Total recoil: the maximum kick from nonspinning black-hole binary inspiral. 2006. eprint gr-qc/0610154.
- [37] T. Goodale, G. Allen, G. Lanfermann, J. Massó, T. Radke, E. Seidel, and J. Shalf. The Cactus framework and toolkit: Design and applications. In *Vector and Parallel Processing – VECPAR’2002, 5th International Conference, Lecture Notes in Computer Science*, Berlin, 2003. Springer.
- [38] F. Herrmann, I. Hinder, D. Shoemaker, P. Laguna, and R. A. Matzner. Gravitational recoil from spinning binary black hole mergers. *Astrophys. J.*, (661):430–436, 2007. eprint gr-qc/0701143.
- [39] F. Herrmann, D. Shoemaker, and P. Laguna. Unequal-mass binary black hole inspirals. *Class. Quantum Grav.*, 24:S33–S42, 2007. eprint gr-qc/0601026.
- [40] S. Husa, I. Hinder, and C. Lechner. Kranc: a Mathematica application to generate numerical codes for tensorial evolution equations. *Comput. Phys. Comm.*, 174:983–1004, 2006. eprint gr-qc/0404023.
- [41] M. Koppitz, D. Pollney, C. Reisswig, L. Rezzolla, J. Thornburg, P. Diener, and E. Schnetter. Recoil velocities from equal-mass binary-black-hole mergers. *Phys. Rev. Lett.*, 99:041102, 27 July 2007. eprint gr-qc/0701163.

- [42] C. Lechner, D. Alic, and S. Husa. From tensor equations to numerical code — computer algebra tools for numerical relativity. In *SYNASC 2004 — 6th International Symposium on Symbolic and Numeric Algorithms for Scientific Computing, Timisoara, Romania, 2004*. eprint cs.SC/0411063.
- [43] N. I. Libeskind, S. Cole, C. S. Frenk, and J. C. Helly. The effect of gravitational recoil on black holes forming in a hierarchical universe. *Mon.Not.Roy.Astron.Soc.*, 368:1381–1391, 2006. eprint arXiv:astro-ph/0512073v1.
- [44] F. Löffler, L. Rezzolla, and M. Ansorg. Numerical evolutions of a black hole-neutron star system in full general relativity: Head-on collision. *Phys. Rev. D*, 74:104018, 2006. gr-qc/0606104.
- [45] P. Madau and E. Quataert. The effect of gravitational-wave recoil on the demography of massive black holes. *Astrophys.J.*, 606:L17–L20, 2004.
- [46] D. Merritt, M. Milosavljevic, M. Favata, S. A. Hughes, and D. E. Holz. Consequences of gravitational radiation recoil. *Astrophys.J.*, 607:L9–L12, 2004. eprint astro-ph/0402057.
- [47] P. Mészáros. Gamma-ray bursts. *Reports of Progress in Physics*, 69:2259, 2006. eprint astro-ph/0605208.
- [48] M. C. Miller and D. P. Hamilton. Four-body effects in globular cluster black hole coalescence. 2002. eprint arXiv:astro-ph/0202298v1.
- [49] R. M. O’Leary, F. A. Rasio, J. M. Fregeau, N. Ivanova, and R. O’Shaughnessy. Binary mergers and growth of black holes in dense star clusters. *Astrophys.J.*, 637:937–951, 2006. eprint arXiv:astro-ph/0508224v2.
- [50] C. D. Ott. *Stellar Iron Core Collapse in 3+1 General Relativity and The Gravitational Wave Signature of Core-Collapse Supernovae*. PhD thesis, Universität Potsdam, Potsdam, Germany, 2006.
- [51] C. D. Ott, A. Burrows, and E. Schnetter. Towards Realistic Models of Long-Soft GRB Central Engines: I. Methods and Polytropic Models. *in preparation*, 2008.
- [52] C. D. Ott, H. Dimmelfmeier, A. Marek, H.-T. Janka, I. Hawke, B. Zink, and E. Schnetter. 3d collapse of rotating stellar iron cores in general relativity including deleptonization and a nuclear equation of state. *Phys. Rev. Lett.*, 98:261101, 2007. eprint astro-ph/0609819.
- [53] C. D. Ott, H. Dimmelfmeier, A. Marek, H.-T. Janka, B. Zink, I. Hawke, and E. Schnetter. Rotating collapse of stellar iron cores in general relativity. *Class. Quantum Grav.*, 24:S139–S154, 2007. eprint astro-ph/0612638.
- [54] A. Peres. Classical radiation recoil. *Phys. Rev.*, 128:2471–2475, 1962.
- [55] D. Pollney, C. Reisswig, L. Rezzolla, B. Szilágyi, M. Ansorg, B. Deris, P. Diener, E. N. Dorband, M. Koppitz, A. Nagar, and E. Schnetter. Recoil velocities from equal-mass binary black-hole mergers: a systematic investigation of spin-orbit aligned configurations. *Phys. Rev. D*, 76:124002, 2007. eprint arXiv:0707.2559 [gr-qc].
- [56] L. Rezzolla, E. Barausse, E. N. Dorband, D. Pollney, C. Reisswig, J. Seiler, and S. Husa. On the final spin from the coalescence of two black holes. arXiv:0712.3541 [gr-qc], 2007, eprint arXiv:0712.3541 [gr-qc].
- [57] L. Rezzolla, P. Diener, E. N. Dorband, D. Pollney, C. Reisswig, E. Schnetter, and J. Seiler. The final spin from coalescence of aligned-spin black-hole binaries. *Astrophys. J.*, 674:L29–L32, Feb. 2008. eprint arXiv:0710.3345 [gr-qc].
- [58] L. Rezzolla, E. N. Dorband, C. Reisswig, P. Diener, D. Pollney, E. Schnetter, and B. Szilágyi. Spin diagrams for equal-mass black-hole binaries with aligned spins. *Astrophys. J. (in press)*, 708, Aug. 2007. eprint arXiv:0708.3999 [gr-qc].
- [59] E. Schnetter, P. Diener, N. Dorband, and M. Tiglio. A multi-block infrastructure for three-dimensional time-dependent numerical relativity. *Class. Quantum Grav.*, 23:S553–S578, 2006. eprint gr-qc/0602104.
- [60] E. Schnetter, S. H. Hawley, and I. Hawke. Evolutions in 3D numerical relativity using fixed mesh refinement. *Class. Quantum Grav.*, 21(6):1465–1488, 21 March 2004. eprint gr-qc/0310042.
- [61] S. E. Woosley and H.-T. Janka. The physics of core-collapse supernovae. *Nature Physics*, 1:147, 2005. eprint astro-ph/0601261.
- [62] B. Zink, E. Schnetter, and M. Tiglio. Multi-patch methods in general relativistic astrophysics – I. hydrodynamical flows on fixed backgrounds. 2007. submitted to PRD, eprint arXiv:0712.0353.