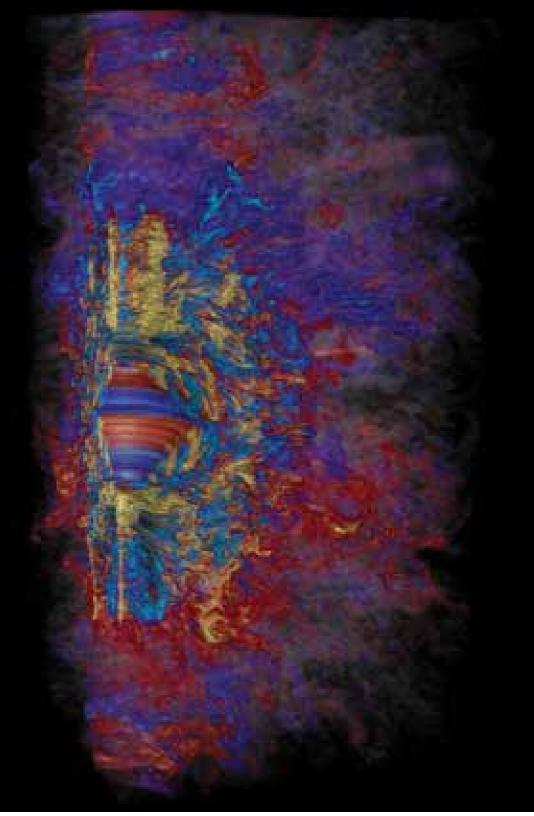

FIGURE 1: Volume rendering of the toroidal magnetic field strength in a rapidly rotating core-collapse supernovae. The simulation shows a forming magnetar as the turbulent engine of the explosion. Colors represent magnetic field strength. Yellow and light blue indicate magnetar-strength positive/negative field and red and blue weaker field.

Image credit: Robert R. Sisneros (NCSA) and Philipp Mösta (UC Berkeley).

MAGNETARS, BLACK HOLE COLLISIONS FOR LIGO, AND A NEXT GENERATION NUMERICAL RELATIVITY CODE

PI: Christian D. Ott[1]
Co-PIs: Mark Scheel[1], Peter Diener[2]
Collaborators: Philipp Mösta[3], Luke Roberts[1], David Radice[1], Sherwood Richers[1], Roland Haas[5], Erik Schnetter[2], Lawrence Kidder[4], Saul Teukolsky[4], Scott Field[4], and Francois Foucart[3]

[1]California Institute of Technology
[2]Lousinana State University
[3]University of California, Berkeley
[4]Cornell University
[5]Albert Einstein Institute

EXECUTIVE SUMMARY

Neutron stars with an extremely strong magnetic field are called magnetars. An open question is how they obtain their magnetic field. We have carried out 3D simulations of a magnetized differentially rotating proto-neutron star. These simulations show that the magnetorotational instability is active in such proto-neutron stars. It efficiently amplifies magnetic field locally and a dynamo process rearranges and orders the field, creating a large-scale global magnetic field as needed for driving supernova explosions.

The National Science Foundation's (NSF) Laser Interferometer Gravitational Wave Observatory (LIGO) has made its first detection of gravitational waves. We responded with target-of-opportunity numerical relativity simulations on Blue Waters that are now helping LIGO to extract information more reliably from the waves.

To more efficiently use large core counts on Blue Waters and in preparation for the next-generation high-performance computing architectures, we have designed, implemented, and tested the new general-relativistic magnetohydrodynamics (MHD) code, SpECTRE. It employs the Discontinuous Galerkin method; is based on the Charm++ parallel programming system, and it exhibits near perfect strong scaling to 128,000 Blue Waters cores.

INTRODUCTION

Numerical general relativity simulations discretize Einstein's equations and are essential to study phenomena involving strong gravity, such as the collision of black holes or the formation of neutron stars in core-collapse supernova explosions.

Magnetars, neutron stars with extremely strong magnetic fields ($B \geq 10^{15}$ Gauss), are believed to power energetic supernova explosions (so-called "hypernovae") and cosmic gamma-ray bursts. MHD turbulence has been suggested as a mechanism for building up the extremely strong magnetar magnetic field. This MHD turbulence is driven by the magnetorotational instability (MRI), whose fastest growing mode has a very short wavelength, making it difficult to resolve in simulations. An additional complication is that the MRI builds up magnetic field strength locally, but what is needed for a magnetar is a large-scale (global) magnetic field that can force the high-density stellar plasma into coherent motion and thus drive a supernova explosion.

Gravitational Waves (GWs) from a coalescing pair of black holes were observed by LIGO on September 14, 2015. This first direct detection of GWs opened up the new field of GW Astronomy. Numerical relativity simulations of coalescing binary systems (and other GW sources) provide essential GW signal predictions for LIGO – these predictions are needed for comparison with GW observations to extract the astrophysical parameters (e.g., mass and spin of black holes) of the sources and for testing General Relativity.

Parallel scaling and efficiency on heterogeneous architectures are key requirements for petascale (and future exascale) simulation codes. Much current computational astrophysics and numerical relativity codes use traditional data-parallel approaches and sequential execution models that suffer from load balancing issues, allow no or only limited execution concurrency and make it difficult to benefit from accelerators. High-order finite-difference and finite-volume methods require the communication of three or more subdomain boundary points in each direction, creating large communication overheads. The Discontinuous Galerkin (DG) method can drastically reduce the number of points that must be communicated. DG, in combination with a task-based parallelization strategy (which allows task concurrency), provides dynamical load balancing, and can hide latencies, may be the route forward for simulation codes.

METHODS & RESULTS

Making Magnetars

We employed our 3D general-relativistic MHD code Zelmani for our magnetar MHD turbulence simulations. Zelmani is open-source, is based on the Cactus framework, and uses the Einstein Toolkit. We mapped our initial conditions from a full 3D adaptive mesh refinement simulation of rapidly rotating stellar collapse with an initial magnetic field of 10^{10} Gauss. We then carried out simulations at four different resolutions (500 m, 200 m, 100 m, 50 m), where the two highest resolutions are sufficient to resolve the key driver of turbulence in this system, the MRI. The 50 m simulation was run on 132,768 Blue Waters cores, used 7.8 TB of runtime memory, and created 500 TB of output data.

Previous local simulations have shown that the instability creates small patches of the magnetar-strength magnetic field, but have not been able to address whether these patches can be turned into a global dynamically relevant toroidal magnetic field structure needed to power an explosion.

Our simulations show that not only does the MRI produce local magnetar-strength magnetic field, but that a subsequent dynamo process connects this local field into a global field that is able to support an explosion. This dynamo process, which is akin to local storm cells merging to form large-scale storm patterns like a hurricane, works extremely efficiently and we predict the formation of a large-scale toroidal field independent of the initial magnetization of the progenitor star. Figure 1 shows an example magnetic field configuration obtained in our highest-resolution simulation.

Our simulations demonstrate that rapidly rotating stellar collapse is a viable formation channel for magnetars. Magnetars are abundantly observed in the universe and have recently gained a lot of attention as possible engines driving hypernovae and superluminous supernovae, some of which are also connected to long gamma-ray bursts.

Binary Black Hole Waveforms for Gravitational Wave Astronomy with LIGO

We simulate binary black hole (BBH) coalescence using the Spectral Einstein Code (SpEC). SpEC uses multi-domain pseudospectral methods to discretize Einstein's equations. Pseudospectral methods provide exponential convergence and are optimal for simulating gravitational fields, which are guaranteed to be smooth. The black hole interiors are excised from the computational grid to avoid singularities.

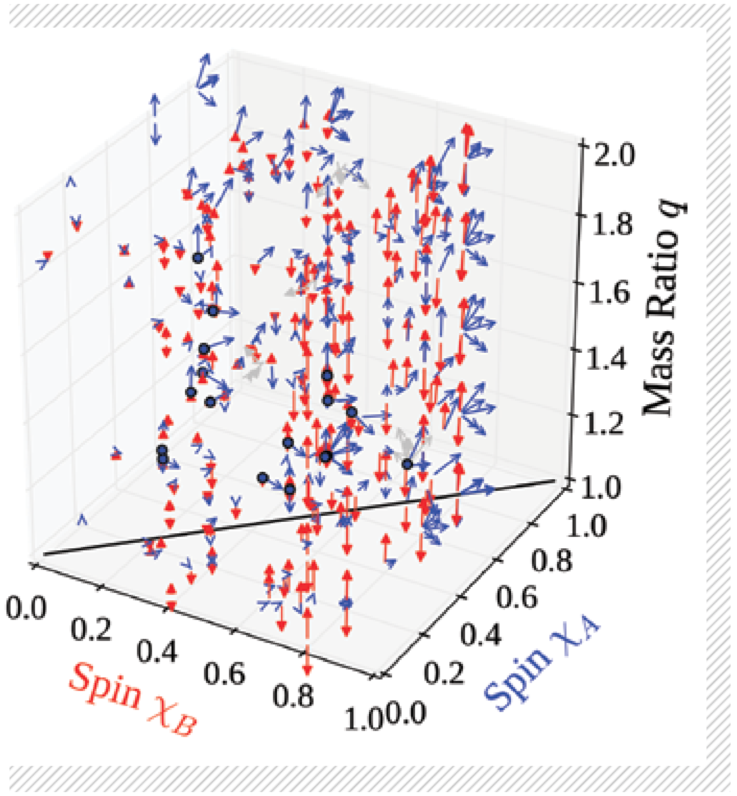

FIGURE 2: Parameters of the 286 binary black hole systems simulated for LIGO. The location of each circle indicates the mass ratio and dimensionless spin magnitudes for a system. The arrows indicate the direction of spin of each black hole (blue arrows for hole A, the larger black hole; red arrows for hole B) and are scaled by the magnitude of the spin.

An astrophysical binary black hole (BBH) system has seven parameters: the mass ratio M_A/M_B and two spin vectors with three components each. The total mass of the system is a scale factor that can be removed from theoretical predictions, but can be inferred from the frequency of the observed waves. Using LIGO data and approximate analytic models for BBH gravitational waveforms, some of the BBH parameters have been measured for the BBH system that LIGO observed. To improve these measurements, numerical relativity waveforms can be used. However, Bayesian parameter estimation algorithms need millions of waveforms as they sample the parameter space near the detected signal. This is problematic because each numerical relativity simulation takes weeks to months to complete.

To address this problem, we have developed a technique to interpolate between numerical relativity waveforms that we *did* simulate to construct waveforms at parameters that we *did not* simulate. The interpolant waveform can be computed in a fraction of a second (compared to the tens to hundreds of thousand central processing unit hours required to construct such a waveform using numerical relativity). Furthermore, the interpolant can be made indistinguishable from the results of numerical relativity to within a specifiable error tolerance. It can thus serve as a *surrogate model* that replaces additional numerical relativity simulations. This error tolerance dictates how many numerical relativity simulations are needed – typically many orders of magnitude fewer than the millions that are needed for parameter estimation – and a greedy algorithm determines where in parameter space these simulations must be performed.

Motivated by LIGO's detection, we have carried out target-of-opportunity SpEC BBH simulations using Blue Waters essentially as a high-throughput capacity computing facility to run hundreds of BBH simulation, each using only two or three nodes. Figure 2 shows the simulated points of the parameter space. In total, we simulated 286 BBH systems, and each was simulated with three different resolutions to estimate the truncation error. Using these waveforms, we built a surrogate model that outperforms previous analytic waveform models and enables more accurate parameter estimation.

A next-generation simulation code: SpECTRE

SpECTRE is a new numerical relativity and general-relativistic hydrodynamics code that uses the DG method with a task-based parallelization strategy facilitated by the Charm++ parallel programming system. DG and task-based parallelization make a promising combination that will allow multi-physics applications to be treated both accurately and efficiently. On Blue Waters, SpECTRE now scales to more than 128,000 cores on a relativistic MHD test case simulating the interactions of shocks (Fig. 3). The Charm++ task-based parallelism library includes a dynamic runtime system to assign task ordering as well as communication optimizations tailored to Cray's Gemini network.

Building on this success, we are now adding the remaining infrastructure and physics to simulate realistic astrophysical systems. To tackle challenging multi-scale problems which naturally arise in such systems, we are adding local time stepping techniques and adaptive mesh refinement strategies to either split the elements into smaller elements or increase the number of basis functions in an element.

WHY BLUE WATERS

Our MHD turbulence simulations that resolve the MRI would have been **completely impossible without Blue Waters'** capability. Blue Waters was essential for these simulations and directly facilitated the breakthrough in our understanding how magnetar-strength magnetic field can be generated. Blue Waters' capacity is crucial for a quick turnaround in our target-of-opportunity numerical relativity BBH simulations. No other machine allows us to run hundreds of simulations concurrently to generate the theoretical predictions needed for LIGO. Finally, Blue Waters is a highly-valuable development and testing resource for our next-generation code SpECTRE. We couldn't develop this code without being able to test it at scale on Blue Waters.

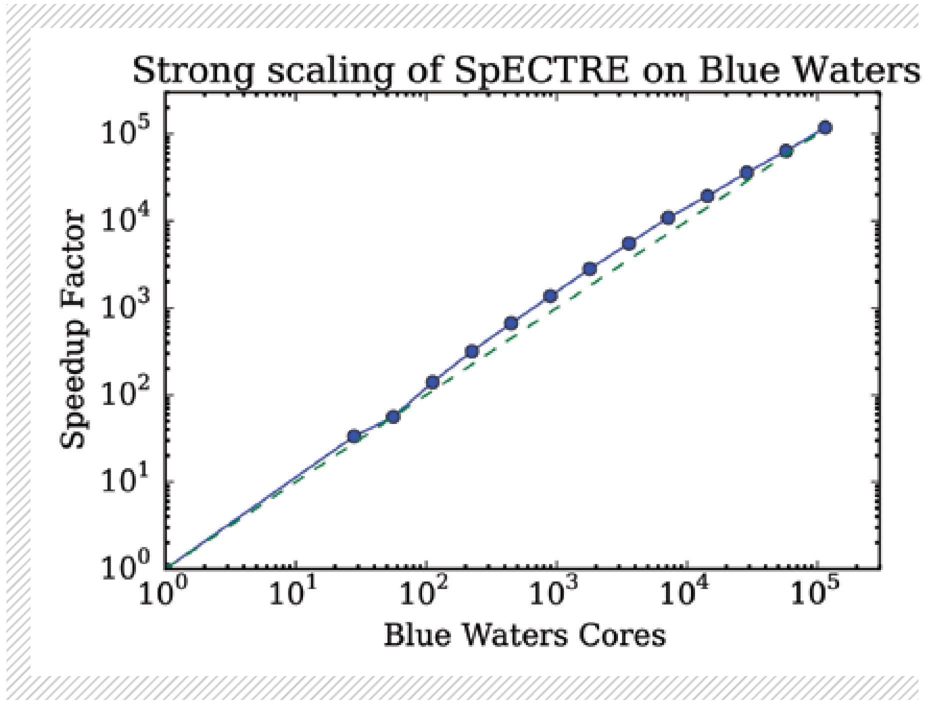

FIGURE 3: Strong scaling of the next-generation general-relativistic MHD code SpECTRE for an MHD test problem on Blue Waters.

NEXT GENERATION WORK

We are working with high intensity on our new code SpECTRE, which we plan to deploy on a next-generation Track-1 system for more detailed, higher-resolution, and longer (in physical time) simulations of core-collapse supernovae, magnetars, and binary systems of black holes and neutron stars. Present simulations, even with Blue Waters, are limited in accuracy and, importantly, in the physical time they can track. Our magnetar MHD simulation, for example, could cover only tens millisecond of the magnetar's life while we would **need to simulate at least a hundred times longer** to fully study its magnetic field evolution and its impact on the supernova explosion.

PUBLICATIONS AND DATA SETS

Mösta P., et al., A large-scale dynamo and magnetoturbulence in rapidly rotating core-collapse supernovae, *Nature*, 528:376 (2015), doi: 10.1038/nature15755

Radice D., et al., Neutrino-driven Convection in Core-collapse Supernovae: High-resolution Simulations, *Astrophys. J.*, 820:76 (2016), doi: 10.3847/0004-637X/820/1/76

Radice D., et al., The One-Armed Spiral Instability in Neutron Star Mergers and its Detectability in Gravitational Waves, *Phys. Rev. D.*, arXiv:1603.05726 (2016).

Radice D., et al., Dynamical Mass Ejection from Binary Neutron Star Mergers. *Mon. Not. Roy. Ast. Soc.*, arXiv:1601.02426 (2016).

Bernuzzi S., et al., How Loud Are Neutron Star Mergers? *Mon. Not. Roy. Ast. Soc.*, arXiv:1512.06397 (2016).