5-2020

A Tale of Two (City-Sized) Nuclei: Investigating Neutron Star Mergers with WhiskyTHC

Zachary Jones Carter
Trinity University, zcarter309@gmail.com

A Tale of Two (City-Sized) Nuclei: Investigating Neutron Star Mergers with WhiskyTHC

by

Zachary J. Carter

An honors thesis submitted to the Department of Physics & Astronomy at

TRINITY UNIVERSITY

in partial fulfillment of the requirements for the
Bachelor of Science in Physics
May 2020

Accepted by .....
Prof. Kelvin Cheng

Accepted by .....
Prof. Nirav Mehta

Accepted by .....
Prof. David Pooley

Accepted by .....
Prof. Orrin Shindell

Accepted by .....
Prof. Jennifer Steele, Chair

Accepted by .....
Prof. Niescja Turner

Accepted by .....
Prof. Dennis Ugolini

A Tale of Two (City-Sized) Nuclei: Investigating Neutron Star Mergers with WhiskyTHC

by

Zachary J. Carter

Submitted to the Department of Physics & Astronomy
on 2020-Apr-29, in partial fulfillment of the
requirements for the Bachelor of Science in Physics

Abstract

The recent success of neutron star merger observations, epitomized by the detection of gravitational and electromagnetic signals from GW170817, necessitates further theoretical work in the area. The program WhiskyTHC stands as one of the most advanced neutron star merger simulators to date. We use the results of WhiskyTHC simulations to establish equation-of-state-dependent relationships between binary mass composition and properties of the produced gravitational radiation, namely total emitted energy and angular momentum.

Thesis Supervisor: Prof. David Pooley

Acknowledgments

I would like to thank David Radice, Luciano Rezzolla, and Filippo Galeazzi for creating the workhorse program of this thesis, WhiskyTHC, with a special thank you to Dr. Radice for patiently answering my slew of emails regarding his work. I would also like to thank Drs. Mehta and Ugolini for granting me the scholarship that allowed me to attend Trinity University in the first place. Finally, I would like to thank Dr. Pooley for gently and passionately guiding me through the world of physics for the past four years and onward into the next chapter of my life.

Contents

Chapter 1

Introduction

Second only to black holes themselves, neutron stars are some of the most extreme objects the cosmos has to offer. Indeed, the mass of two Suns squeezed into a ball the size of a large city is bound to boast a bounty of fascinating physical phenomena. However, as is often the case in nature, with fascination comes complexity. This complexity, despite entertaining challenge-hungry scientists over the past several decades and even drawing new minds into the field, has hindered the scientific community's collective understanding of these astronomical oddities. The equation of state, for example, while well understood for ordinary stars (as has been the case for some time now), is still quite uncertain for neutron stars. Indeed, a neutron star is, in essence, an "oversized liquid atomic nucleus" (Bradt 2008) and at the moment, we are not aware of an atomic nucleus's equation of state; in fact, studying neutron stars is one of the most promising ways of determining the nuclear equation of state (Lattimer 2012). Furthermore, the addition of a second neutron star in the case of neutron star mergers (illustrated in Figure 1-1) only magnifies these holes in our understanding. Thus, further investigation of these bodies and their mergers seems prudent.

1.1 Gravitational Waves

The mergers of neutron stars can be detected via electromagnetic as well as gravitational radiation, although, at least in terms of immediate emission, the latter is more likely to be detectable for a larger range of distances (Metzger 2017). Thus, in this sense, gravitational waves can be regarded as the more reliable of the two merger indicators. This precedence, along with gravitational radiation's relative esotericism, marks gravitational waves as a topic worthy of further discussion.

A gravitational wave can most succinctly be described as "a tidal acceleration that propagates" (Taylor, Wheeler, & Bertschinger 2019), or equivalently, a curvature in spacetime that propagates. This propagating curvature makes its presence known by modifying the space through which it passes. To be specific, gravitational radiation stretches space in one direction while compressing space in the orthogonal direction, and vice versa as the wave progresses (Figure 1-2). While gravitational radiation was not formally predicted until 1916, after Einstein had formulated the field equations of general relativity (Abbott et al. 2016), its existence was essentially necessitated upon the formulation of special relativity several years prior.

Indeed, during the era of Newtonian infallibility (that is, the span of time between the publication of Newton's law of universal gravitation in the *Principia* and Einstein's special relativity in *On the Electro-*

Figure 1-1: Artist's depiction of two neutron stars colliding. Credit: NASA/Swift/Dana Berry

dynamics of Moving Bodies), scientists assumed that gravitational information traveled instantaneously. If a giant planet were to suddenly materialize near Earth, the Newtonian scientists of this time would claim that the Earth would experience this new body's gravitational pull instantly, simultaneous with the planet's creation.

However, Einstein shattered such beliefs. His theory of special relativity, which was later experimentally verified, incontrovertibly held that no information can travel faster than light. Everything in the cosmos was held to this universal speed limit, even information. Thus, gravitational information cannot travel infinitely quickly; it must travel at some finite velocity, namely the speed of light. These conclusions were solidified in 1916 once Einstein discovered that in the weak-field limit, his field equations had wave solutions, waves that did in fact travel at the speed of light (Abbott et al. 2016).

1.1.1 Gravitational Wave Metric and Wave Equation

The distortion of spacetime by gravitational waves can be represented by a transformation of the flat, Minkowski spacetime metric $d\tau^2 = dt^2 - dx^2 - dy^2 - dz^2$ into the gravitational wave metric:

$$d\tau^2 = dt^2 - (1+h)dx^2 - (1-h)dy^2 - dz^2 \tag{1.1}$$

where we have assumed that the gravitational radiation is of the form of a plane wave moving in the positive z-direction. This plane wave assumption is a highly realistic one for detection purposes, considering the only sources powerful enough to produce detectable gravitational waves are located extremely far away. The quantity h is twice the value of the *strain*, a concept explained below. This global metric (1.1) can be massaged into a more practical, local form by recalling that anything traveling at light speed, such as gravitational radiation, has a "wristwatch time" of $d\tau = 0$, by approximating the differentials as finite deltas, and by rewriting $(1+h)\Delta x^2 - (1-h)\Delta y^2$ as $[(1+h)^{1/2}\Delta x]^2 - [(1-h)^{1/2}\Delta y]^2$ before taking advantage

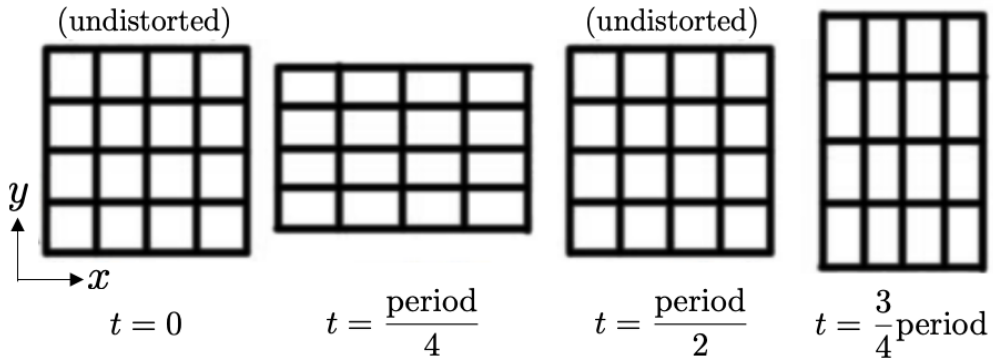

Figure 1-2: Cross-section of space through which a gravitational wave is passing in the positive z-direction, that is, out of the page. As a single wavelength passes through the cross-section of space, the space remains undistorted at two times ($t = 0$ and $t = \frac{\text{period}}{2}$) and the space is altered via a trough or peak of the wave at two times ($t = \frac{\text{period}}{4}$ and $t = \frac{3}{4}$period). At $t = \frac{\text{period}}{4}$, space is stretched in the x-direction and compressed in the y-direction; however, at $t = \frac{3}{4}$period the opposite occurs, as space is compressed in the x-direction and stretched in the y-direction. Credit: Taylor, Wheeler, & Bertschinger (2019)

of the small size of h relative to 1 to simplify the previous expression via the binomial approximation. The result is the following expression:

$$0 = \Delta t_{\text{obs}}^2 - [(1 + \frac{h}{2})\Delta x_{\text{obs}}]^2 - [(1 - \frac{h}{2})\Delta y_{\text{obs}}]^2 - \Delta z_{\text{obs}}^2 \tag{1.2}$$

where the 'obs' subscript is used in order to emphasize that this metric codifies what one would actually observe during the passing of a gravitational wave and where the strain $\frac{h}{2}$ is more explicitly present.

The strain, or gravitational wave strain, describes the fractional change in length experienced by space in the x and y directions. Indeed, given two masses separated by a distance L in the x or y direction, the wave will induce a change in separation of $\frac{h}{2} \times L$. It is this change in separation that enables gravitational wave detection, as the Laser Interferometer Gravitational-Wave Observatory (LIGO) is fine-tuned to detect these comically minuscule deviations in length along its legs (Taylor, Wheeler, & Bertschinger 2019). Thus, the gravitational wave strain $\frac{h}{2}$ is the key observable in gravitational wave observations of neutron stars. The details of this observation process are laid out in Section 1.1.2 below. The value of h ultimately stems from Einstein's field equations, which actually reduce to a wave equation for flat spacetime and small h values, that is, the conditions we would expect here on Earth. The wave equation can be most generally expressed as

$$\frac{\partial^2 h}{\partial t^2} = \frac{\partial^2 h}{\partial x^2} + \frac{\partial^2 h}{\partial y^2} + \frac{\partial^2 h}{\partial z^2} \tag{1.3}$$

which we can solve for a plane wave propagating in the positive z-direction (Taylor, Wheeler, & Bertschinger 2019).

1.1.2 Detection

On September 14, 2015, LIGO made its first official gravitational wave detection. The signal was thus dubbed GW150914 after the year, month, and day of the reception. The waves had been produced by a

collision of two black holes 230 to 570 Mpc away from Earth (Abbott et al. 2016). In order to understand the mechanics of this revolutionary detection, we must again look towards the metric (1.1).

In order to explain measurements made here on Earth, it would be useful to adapt (1.1) to a local inertial frame, namely the Earth frame. First, we should establish the general form of such a local metric:

$$\Delta\tau^2 \approx \Delta t_{\text{Earth}}^2 - \Delta x_{\text{Earth}}^2 - \Delta y_{\text{Earth}}^2 - \Delta z_{\text{Earth}}^2 \tag{1.4}$$

Comparing (1.4) with the approximate form of (1.1):

$$\Delta\tau^2 \approx \Delta t^2 - (1+h)\Delta x^2 - (1-h)\Delta y^2 - \Delta z^2 \tag{1.5}$$

we may conclude that

$$\Delta t_{\text{Earth}} = \Delta t \tag{1.6}$$

$$\Delta x_{\text{Earth}} = (1+h)^{1/2}\Delta x \approx (1+\frac{h}{2})\Delta x \tag{1.7}$$

$$\Delta y_{\text{Earth}} = (1-h)^{1/2}\Delta y \approx (1-\frac{h}{2})\Delta y \tag{1.8}$$

$$\Delta z_{\text{Earth}} = \Delta z \tag{1.9}$$

where a binomial approximation is utilized in (1.7) and (1.8), as $h \ll 1$ by the time the wave gets to Earth.

Now, (1.6) and (1.9) insist that time as well as space along the direction of gravitational wave propagation are the same for map (global, arbitrary coordinates) and Earth coordinates. More importantly for detection purposes however, (1.7) and (1.8) clearly imply that the x and y intervals differ. Relating this math to LIGO, note that for a stationary mass, like one of LIGO's suspended mirrors, the map Δx and Δy will remain unchanged by the Principle of Maximal Aging. However, as the gravitational wave passes by and h changes, Δx_{Earth} and Δy_{Earth} will change. To elaborate, say for $h > 0$, rest masses separated in the x-direction will move away from each other while rest masses separated in the y-direction will move closer together, and vice versa as h oscillates. As a result, there is alternating stretching and compression in perpendicular directions!

Thus, as a gravitational wave passes through LIGO, unless the wave's polarization is oriented exactly 45° from the interferometer legs (thereby causing both legs to be identically influenced by the wave at a given time) the laser beam has a different distance to travel depending on the leg through which it is traveling. However, by special relativity, light will still travel at the same speed in the local Earth frame. Therefore, considering the laser beams now have two different distances to travel while maintaining the same speed, they must experience different travel times by the classic $t = d/v$ where t is time elapsed, d is distance traveled, and v is velocity. This time difference Δt_{Earth} can be easily calculated. Holding D to be the measured leg lengths with no gravitational radiation present (and thus $D = \Delta x = \Delta y$ within the context of (1.7) and (1.8)), we can infer from subtracting (1.8) from (1.7), and subsequently applying $t = d/v$ (holding the speed of light to be 1) that the time difference for a single round trip of light is

$$\Delta t_{\text{Earth}} = 2Dh \tag{1.10}$$

assuming the polarization of the gravitational waves lines up perfectly with the interferometer legs. However, light is reflected many times along a LIGO leg; thus, the total time difference is

$$\Delta t_{\text{Earth}} = 2NDh \tag{1.11}$$

where N is the total number of round trips taken by a light beam before undergoing interference.

Keeping in mind that for LIGO $N = 300$ and $D = 4$ km (Taylor, Wheeler, & Bertschinger 2019) and that the max strain $\frac{h}{2}$ of GW150914 was on the order of 10^{-21} (Abbott et al. 2016)), (1.11) reveals that the maximum time difference was roughly 4.8×10^{-15} light-meters or, in more traditional units, 1.6×10^{-23} seconds, a comically small amount of time! In terms of phase shift, there was evidently a shift of 4.8×10^{-15} m between the two laser beams. Considering LIGO utilizes a laser of wavelength 1064 nm (Abbott et al. 2016), we see that the total phase shift amounted to a shift of 4.5×10^{-9} wavelengths. This shift is clearly extremely small; however, it evidently produced a change in the interference pattern significant enough to be detected by LIGO and confirm to the world that gravitational waves exist. A figure qualitatively illustrating the detection process is included for clarification (Figure 1-3).

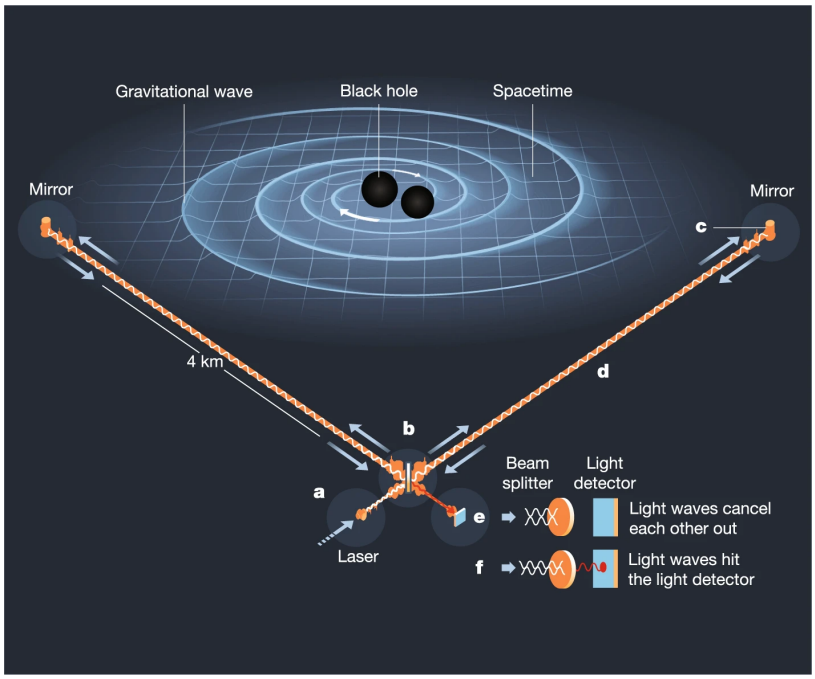

Figure 1-3: A schematic of LIGO's detection system coupled with an image of an inspiraling black hole binary producing gravitational waves. Laser light enters the detector at **a** and is split via a beam splitter at **b**. The beams are repeatedly reflected between **c** and **b** such that each beam makes 300 round trips across the length of its respective leg (**d**). Due to the relative change in leg length induced by passing gravitational waves, the interference pattern projected onto the light detector (**e**) changes with time (**f**). Note that the phase shift, and consequential change in interference pattern, illustrated between **e** and **f** is highly exaggerated. Credit: Miller & Yunes (2019)

1.1.3 Sources

In a sense, gravitational wave production can be somewhat tricky. By Birkhoff's theorem, a spherically symmetric collapse or explosion cannot generate gravitational radiation (or even electromagnetic radiation for that matter). Indeed, it turns out that only motion that lacks spherical and linear symmetry can generate gravitational radiation. Examples of such asymmetric motion frequently occur at the end of a star's life in the form of supernovae explosions as well as subsequent collapses into neutron stars and black holes; both actions are almost certainly asymmetric in reality, thereby avoiding the prohibition imposed by Birkhoff.

Binary orbits are also suitably asymmetric, yet they are also simple enough to earn the distinction of being the only gravitational wave source for which the wave equation can be explicitly solved:

$$h(z,t) = -\frac{4G^2 M_1 M_2}{c^4 r z} \cos\left[\frac{2\pi f(z-ct)}{c}\right] \tag{1.12}$$

where h is again twice the strain, M_1 and M_2 are the component masses, r is their orbital separation, z is the distance between the binary and the observer, and f is the gravitational wave frequency, which is simply twice the orbital frequency f_{orbit}, a simple Newtonian calculation:

$$f_{\text{orbit}} = \frac{1}{2\pi}\left[\frac{M_1 + M_2}{r^3}\right]^{1/2}. \tag{1.13}$$

Note that this equation assumes that the binary lies in the x-y plane (Taylor, Wheeler, & Bertschinger 2019).

Thus, our theoretical understanding of binary systems' gravitational radiation is quite firm. Considering a merger event is just the coalescence of the members of a binary, this understanding has been of invaluable use to the studies of mergers and their constituents and has flagged merger events as a particularly desirable target of gravitational wave observations. Indeed, considering mergers are enabled by orbital energy loss primarily accomplished via gravitational radiation (Taylor, Wheeler, & Bertschinger 2019), their importance in gravitational wave astronomy is quite natural.

1.1.4 Discernible Information

Gravitational radiation can provide a bounty of information regarding its source. In the case of mergers, the frequency and amplitude of the pre-coalescence waveform allow one to constrain the component masses of a merger, while the remnant mass can be ascertained from the frequency of the post-coalescence waveform, at least for the case of a black hole remnant. Also, the orientation of the wave, which can be deduced by comparing the signal's temporal offsets between observatories, helps determine in what direction the source is located (Taylor, Wheeler, & Bertschinger 2019). More specifics regarding the information extractable from neutron star-neutron star merger signals are discussed in the case study below.

1.2 GW170817 Case Study

Gravitational wave investigations of neutron star mergers are best handled by LIGO, as evidenced by their revolutionary detection of GW170817. This signal, produced by colliding neutron stars of comparable masses (both lying between 1.1 and 1.6 $M_\odot$), stands as the first ever detection of neutron star-neutron star merger gravitational waves and, arguably more importantly, as a milestone of multimessenger astronomy. Indeed, both gravitational and electromagnetic waves were detected from the merger, marking the first time both kinds of radiation were observed from the same astronomical event (Metzger 2017).

From the gravitational signal alone astronomers were able to ascertain the radii of the colliding neutron stars ($\lesssim$ 13 km). The effective tidal deformablility of the system, a parameter directly measurable from the gravitational waveform, had previously been thought to only lend itself to a calculation of compactness (a function of mass as well as radius). However, Raithel et al. (2018) determined that in the case of neutron star-neutron star mergers, the effective tidal deformability is actually virtually independent of the component masses and instead depends more strongly on the ratio of the chirp mass (another parameter that can be extracted from the waveform) to the component radius. Thus, given an accurate measurement of the chirp mass, one can utilize the effective tidal deformability as an effective, direct probe of neutron star radius. In this manner neutron stars' radii can be calculated without reference to their masses, masses which are typically estimated via several, potentially unreliable, assumptions regarding admissible values of the neutron stars' spins (Abbott et al. 2017).

Even more information was revealed to scientists upon investigation of the accompanying electromagnetic signals. The initial light indicated the incidence of a gamma ray burst (illustrated in Figure 1-4), dubbed GRB170817A. However, perhaps more significant than this signal's identification as a gamma ray burst was its duration: both components were less than or equal to a few seconds. Thus, the merger produced a "short" GRB, a type of gamma ray burst previously hypothesized to be produced by neutron star-neutron star mergers but never empirically linked with these cosmic collisions until the arrival of GW170817 and GRB170817A (Metzger 2017).

Later electromagnetic signals revealed the presence of a markedly distinct astrophysical phenomenon: a kilonova (illustrated in Figure 1-5). Over six decades ago, Burbidge et al. (1957) and Cameron (1957) realized that about half of the elements heavier than iron must be created in an extremely neutron-rich environment, wherein neutron-neutron capture occurs much more quickly than does β-decay. The natural progenitor of this "r-process" was unknown; however, theoretical work eventually propped up neutron star-neutron star mergers as the primary candidates. If correct, this hypothesis demanded that light signatures from a cloud of radioactive decay, namely a kilonova, should be detected in conjunction with a neutron star merger. Indeed, about one day after GW170817, astronomers did begin to detect such decay signatures in the near-infrared and optical, thereby confirming long-standing suspicions regarding the origin of several elements.

Furthermore, through analysis of both the gravitational and electromagnetic signatures of GW170817, the majority of interested scientists seem to have come to the (debated) conclusion that the final remnant of the merger was a black hole, rather than another neutron star (Metzger 2017). In this manner, multimessenger observation of the GW170817 merger has provided empirical support to past theoretical work suggesting that black holes are the typical result of neutron star-neutron star mergers (Fryer et al.

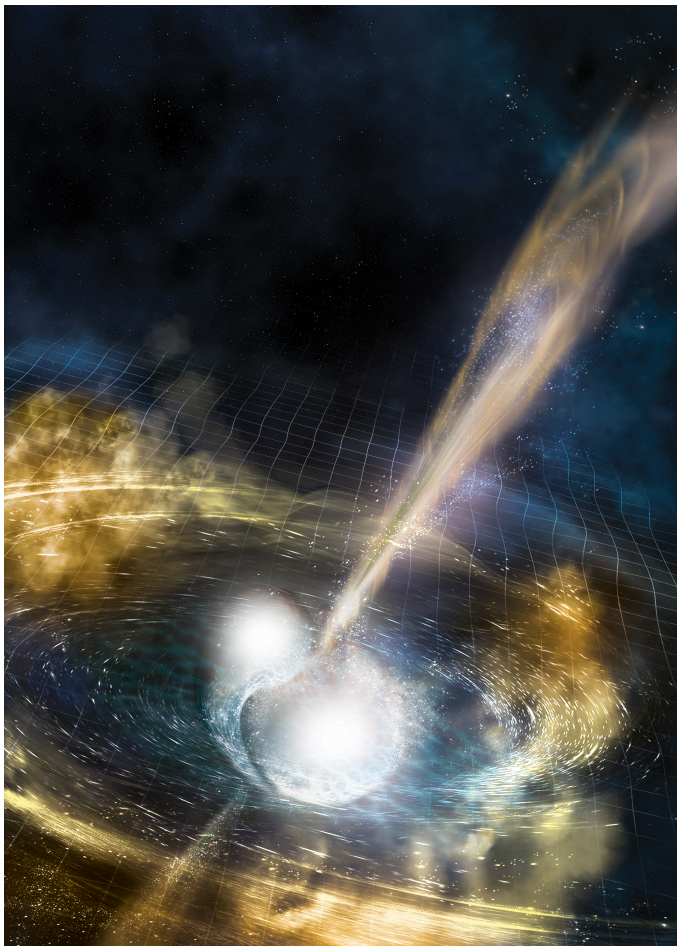

Figure 1-4: Artist's rendition of neutron stars merging with a gamma ray burst firing along the polar axis. A grid is included to illustrate the associated gravitational waves' distortion of spacetime. Credit: National Science Foundation/LIGO/Sonoma State University/A. Simonnet

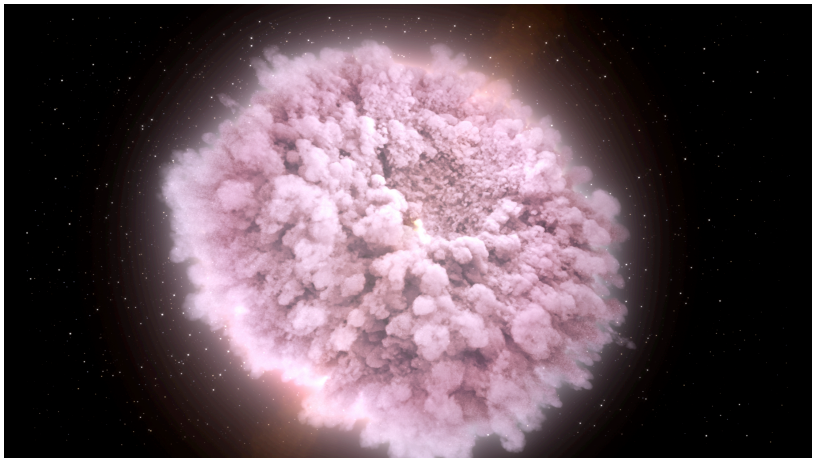

Figure 1-5: Illustration of a kilonova produced by a neutron merger. Credit: NASA's Goddard Space Flight Center/CI Lab

2015). Indeed, if this first round of results is any indication, the multimessenger era promises to accelerate neutron star studies to an unprecedented pace and, to quote Metzger (2017), produce a "quantum leap" in our understanding.

1.3 Theoretical Demand

With the recent, extraordinary success of observational research, theoretical studies of neutron stars and their mergers have become all the more necessary. After all, only through greater theoretical understanding of these mergers can we hope to provide observations with physical significance to any satisfactory detail. Such theoretical examinations are best handled by a program capable of holistically modeling the highly advanced physics, most notably fluid dynamics, general relativity, and neutrino transport, inherent to neutron stars and their collisions, a program like WhiskyTHC.

1.3.1 WhiskyTHC

WhiskyTHC is a high-order, fully general-relativistic hydrodynamics program; that is, as implied above, it can accurately and simultaneously process both the relativity and the fluid mechanics inherent to a given event. Developed by David Radice, Luciano Rezzolla, and Filippo Galeazzi, it is the first code to have achieved "higher than second-order convergence in the calculation of the gravitational radiation from inspiraling binary neutron stars", enabling more intensive studies of the gravitational waves emitted by neutron star binaries.

Codes emblematic of the previous state-of-the-art were unable to consistently produce high-quality waveforms over the parameter space (binary configuration, neutron star compactness, and equation of state) of interest. They were limited by their relatively low convergence order (approximately ≤ 2). By raising the convergence order, WhiskyTHC constitutes a significant step forward in terms of computational accuracy; its gravitational wave phase error, for instance, is about 50 times smaller than that of Whisky, a prime example of the previous state-of-the-art (Radice et al. 2014).

By running simulations over a variety of merger parameters, such as those listed above, I can draw relationships between such parameters and the results of the merger. In this manner, I hope to further our understanding of neutron star mergers and, more fundamentally, of neutron stars themselves.

Chapter 2

Procedure & Results

2.1 Change in Plans

In spite of our best efforts and intentions, we were unable to run WhiskyTHC on Trinity University's computing cluster or at the University of Texas at Austin's Texas Advanced Computing Center (TACC). Indeed, we spent months trying to get the code to compile on Trinity's cluster before learning from a TACC colleague that the program was actually quite fragile and would certainly present us with a new host of problems even after finally attaining compilation. An instance of email correspondence with this TACC member is included in the Appendix, Section 4.1.

Thus, in the interest of time, we abandoned our original plan to conduct these simulations ourselves; instead, we will be taking advantage of past simulations. Drawing from published WhiskyTHC data, we will reproduce a result from the associated literature before exploring our own novel relationships. In particular we will first reproduce Figure 3 of Bernuzzi et al. (2016), which graphs the total normalized gravitational wave energy emitted by various neutron star binaries (differing by the component masses and the equations of state utilized) up to 20 ms after merger (defined by Bernuzzi et al. (2015) as occurring when the amplitude of the $l = 2$, $m = 2$ mode peaks) as a function of the tidal coupling constant. We will then proceed to plot the total emitted energies and the final angular momenta of a sample of binaries as functions of the total masses and component mass ratios.

2.2 Available Data

As implied above, we decided to draw from the data published alongside Bernuzzi et al. (2016). The data contains information regarding the gravitational radiation produced in 24 neutron star merger simulations. Each simulation constitutes a unique combination of one of four mass configurations (in solar masses: $M_1 = 1.40$ and $M_2 = 1.20$, $M_1 = 1.365$ and $M_2 = 1.25$, $M_1 = 1.35$ and $M_2 = 1.35$, and $M_1 = 1.44$ and $M_2 = 1.39$), one of two spatial resolutions ($\Delta x = 295$ m and $\Delta x = 185$ m), and one of three equations of state (in descending stiffness: DD2, LS220, and SFHo). These specific equations of state were chosen by Bernuzzi et al. (2016) in order to "span a reasonable range of radii and maximum gravitational masses for non-spinning" neutron stars and thereby attempt to cover the breadth of physically feasible results of these mergers. DD2, LS220, and SFHo induce radii (of 1.35 $M_\odot$ neutron stars) of 13.2, 12.7, and

11.9 km respectively and maximum masses of 2.42, 2.05, and 2.05 $M_\odot$ respectively. Different simulations are referred to via listing their equation of state followed by their mass configuration and an indication of their resolution via a label of "HR" for the higher resolution ($\Delta x = 185$ m) simulations and a lack of additional notation for the standard resolution ($\Delta x = 295$ m) simulations. For example, a standard resolution simulation of a binary that features individual masses $M_1 = 1.44$ and $M_2 = 1.39$ and utilizes the DD2 equation of state is labeled as DD2-144139.

To illustrate the composition of the data produced during a given simulation, we plot some of the outputs of the aforementioned simulation DD2-144139 below. We can examine how properties of the radiation evolve over time, properties such as amplitude (Figure 2-1), frequency (Figure 2-2), and phase (Figure 2-3). We can also study the total energy emitted via the gravitational waves up to a given point in time (Figure 2-4) as well as the binary angular momentum remaining at a given time as angular momentum is carried off via these same gravitational waves (Figure 2-5).

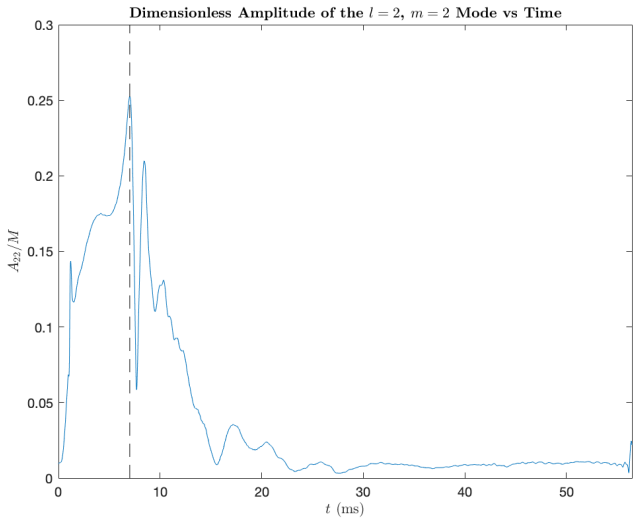

Figure 2-1: Amplitude of the gravitational radiation's $l = 2$, $m = 2$ mode at a distance of 1 $M_\odot$, made dimensionless via division by the binary mass, plotted against time for DD2-144139. The black dotted line marks the time of merger. Note how, as guaranteed by definition, the amplitude peaks at this time.

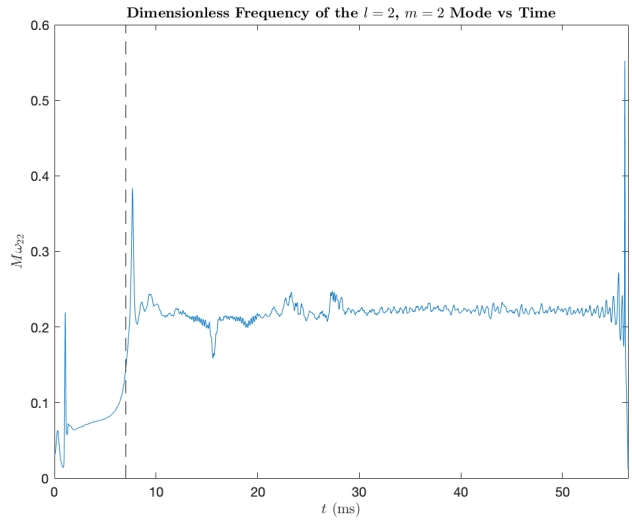

Figure 2-2: Frequency of the gravitational radiation's $l = 2$, $m = 2$ mode in solar masses, made dimensionless via multiplication by the binary mass, plotted against time for DD2-144139. The black dotted line marks the time of merger. Note the extreme rise in frequency at this time. This rapid frequency escalation is the "chirp" characteristic of mergers' gravitational signals.

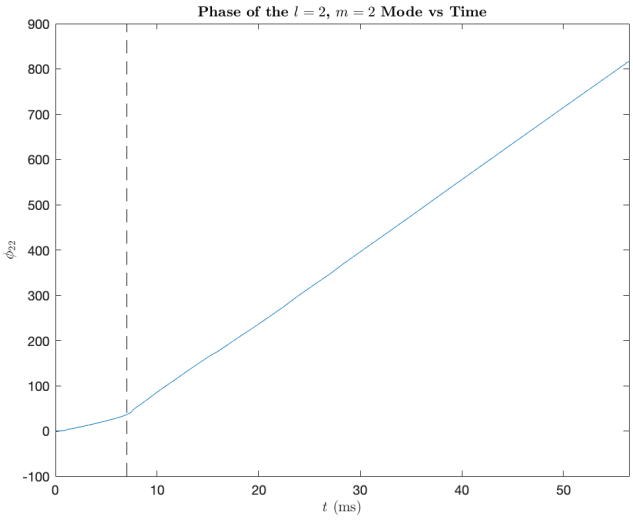

Figure 2-3: Phase of the gravitational radiation's $l = 2$, $m = 2$ mode plotted against time for DD2-144139. The black dotted line marks the time of merger. Note the change in slope at this time.

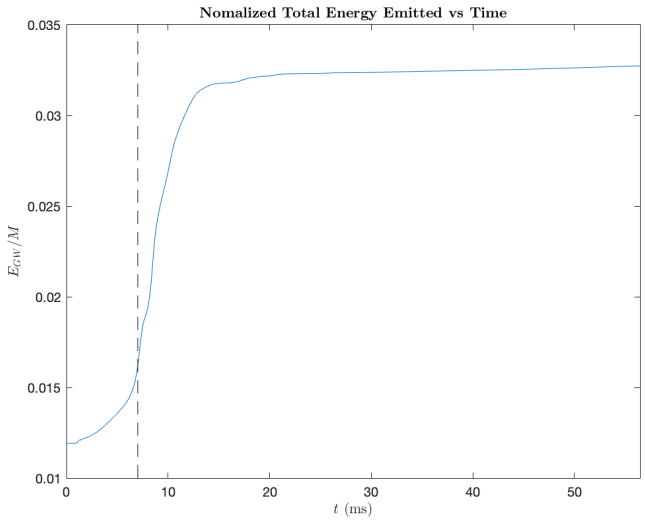

Figure 2-4: Total energy in solar masses emitted via gravitational waves by the time t, normalized via division by the binary mass, plotted against time. The energy plotted includes energy emitted during the inspiral before the simulation began. This initial emitted energy is determined by subtracting the binary mass calculated at the beginning of the simulation $M_{ADM} = 2.799$ $M_\odot$ from the binary mass at infinite separation $M = M_1 + M_2 = 2.83$ $M_\odot$. The black dotted line marks the time of merger. Note how the rate of energy emission seems to peak around this time.

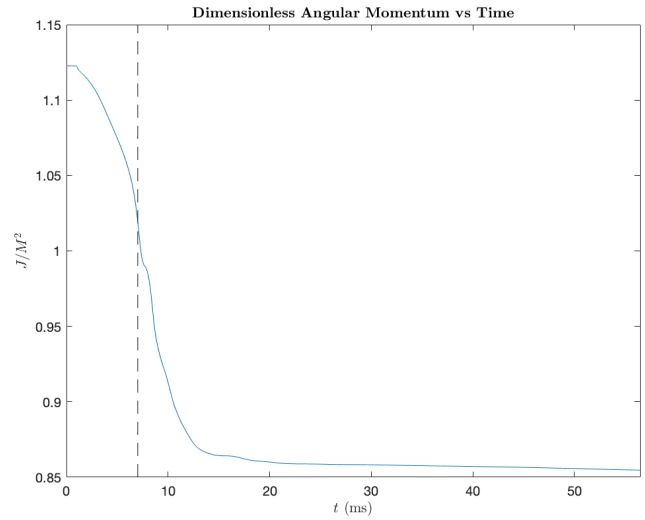

Figure 2-5: Angular momentum of the binary in $M_\odot^2$, made dimensionless via division by the binary mass squared, plotted against time. The angular momentum at a given time t is calculated by subtracting the angular momentum radiated away since the beginning of the simulation up to t from the angular momentum calculated at the beginning of the simulation $J_{ADM} = 7.589$ $M_\odot^2$. The black dotted line marks the time of merger. Note how the angular momentum seems to change most rapidly around this time.

2.3 Figure 3 of Bernuzzi et al. (2016)

Using the data set of Bernuzzi et al. (2016), we ventured to reproduce one of their cornerstone figures, their Figure 3, in order to ensure that our understanding of the data set was sufficient to produce our own novel results. In this figure, for each simulation, the normalized total energy emitted up to 20 ms after the merger (at time t_{20}) is plotted against the tidal coupling constant, an indication of the degree of tidal interaction. If the merger remnant collapsed into a black hole before t_{20}, the total energy emitted up to 1 ms after the collapse time was plotted instead. The figure shows a somewhat robust (as indicated by the relative alignment of the standard and high resolution simulations' data points) relationship between the tidal coupling constant and the emitted energy. The authors thus suggest that, given the tidal coupling constant's correspondence with the "stiffness" of the equation of state, observational constraints could be placed on the actual neutron star equation of state via gravitational wave energy measurements. In light of this information, we consider Figure 3 to be the central result of Bernuzzi et al. (2016), hence our particular efforts to reproduce it.

The fruit of this endeavor is displayed in Figure 2-6 below. The most prominent difference between our recreation and the original is the lack of several points. These points represented those simulations in which the merger remnant collapsed before t_{20}. The exact collapse times of these simulations are no longer available (D. Radice, private communication); thus, calculating the energy emitted up to 1 ms after the collapse time is quite difficult. For this unfortunate reason, we elected to discard these points. However, the remaining points appear to match those of the original figure quite well. Most seem to lie within about 2×10^{-4} of each other in terms of the E_{GW}/M axis with none seeming to stray farther than 4×10^{-4} apart. Although the positive correlation between E_{GW}/M and κ_T^2 at smaller κ_T^2 is absent in our figure due to the lack of the aforementioned early-collapsing points, the negative correlation at higher κ_T^2 is well-maintained. Overall, our reproduction of Bernuzzi et al. (2016)'s Figure 3 is relatively faithful and demonstrates a level of competency with the data sufficient for the purposes of producing our own findings.

2.4 Our Results

Using the data set described above, for each standard resolution simulation, we plotted the total energy radiated by the end of the simulation against the binary mass (Figure 2-7) and the component mass ratio (Figure 2-8). Additionally, we plotted the binary angular momentum at the end of the simulation against, again, the binary mass (Figure 2-9) and the component mass ratio (Figure 2-10). We include high resolution data in the figures displayed in Section 4.2 of the Appendix. We discuss the implications of these plots in the following chapter.

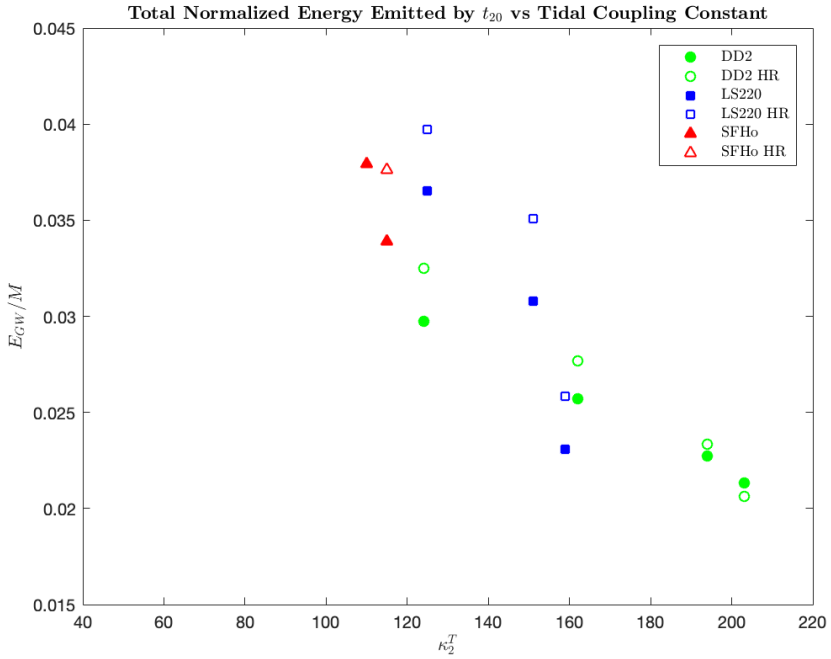

Figure 2-6: Figure 3 of Bernuzzi et al. (2016). The total energy emitted by a binary via gravitational waves up to 20 ms after its merger E_{GW}, normalized by the mass of the binary M, plotted against the binary's tidal coupling constant κ_2^T. The initial radiated energy is calculated in the same manner as described in Figure 2-4 above. The tidal coupling constant κ_2^T describes the severity of the tidal interactions experienced by the binary and is listed for each simulation within the text of Bernuzzi et al. (2016). Unlike in the original figure, binaries that collapsed into black holes before 20 ms had passed since merger are not included, as the collapse time data for these runs is no longer available.

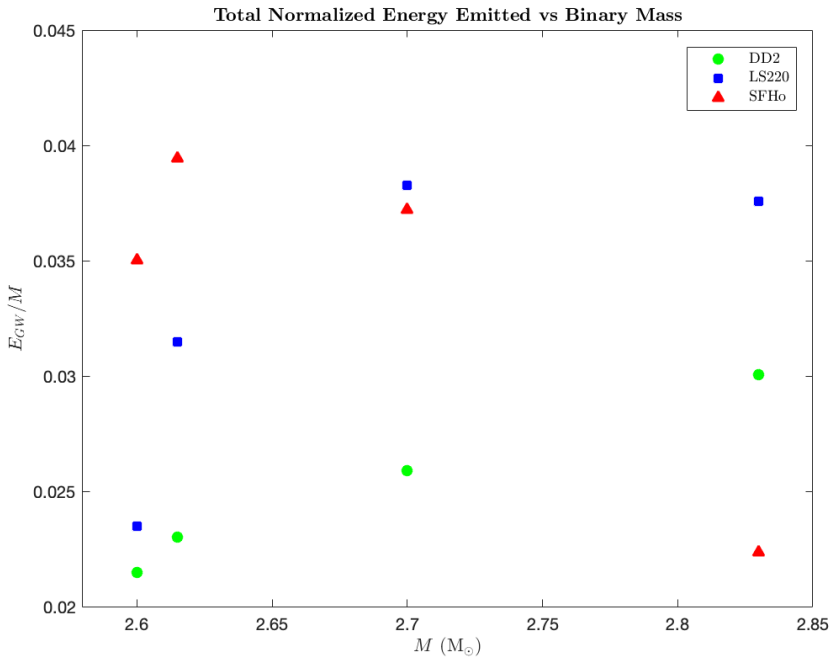

Figure 2-7: The total normalized energy emitted by a binary via gravitational waves by the end of the simulation, including energy radiated before the simulation began, plotted against the binary mass.

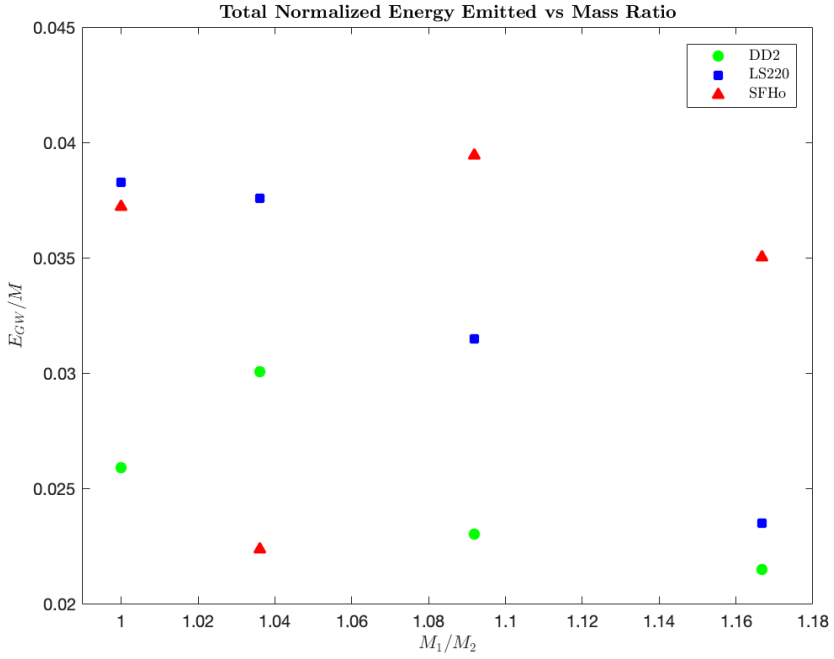

Figure 2-8: The total normalized energy emitted by a binary via gravitational waves by the end of the simulation, including energy radiated before the simulation began, plotted against the mass ratio of the components.

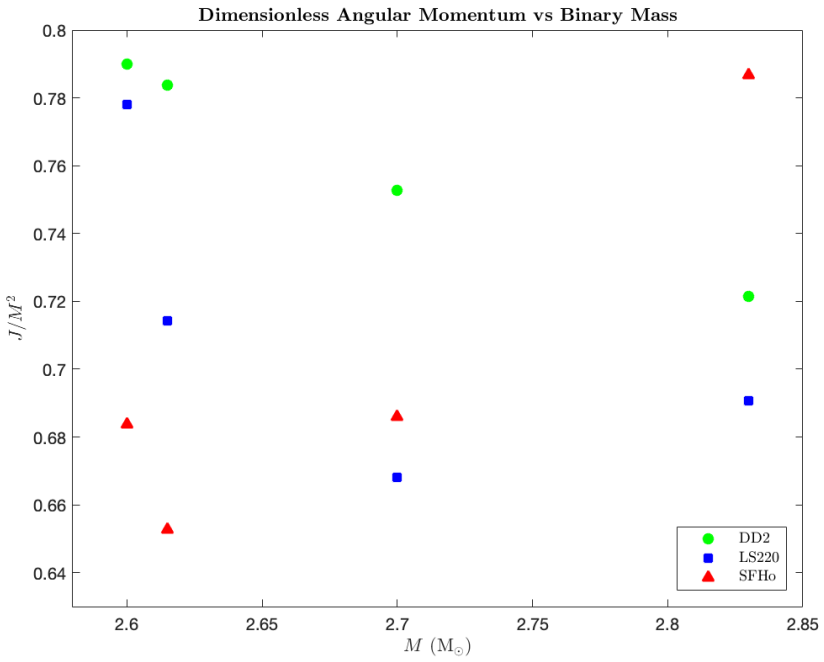

Figure 2-9: The dimensionless binary angular momentum by the end of the simulation plotted against the binary mass.

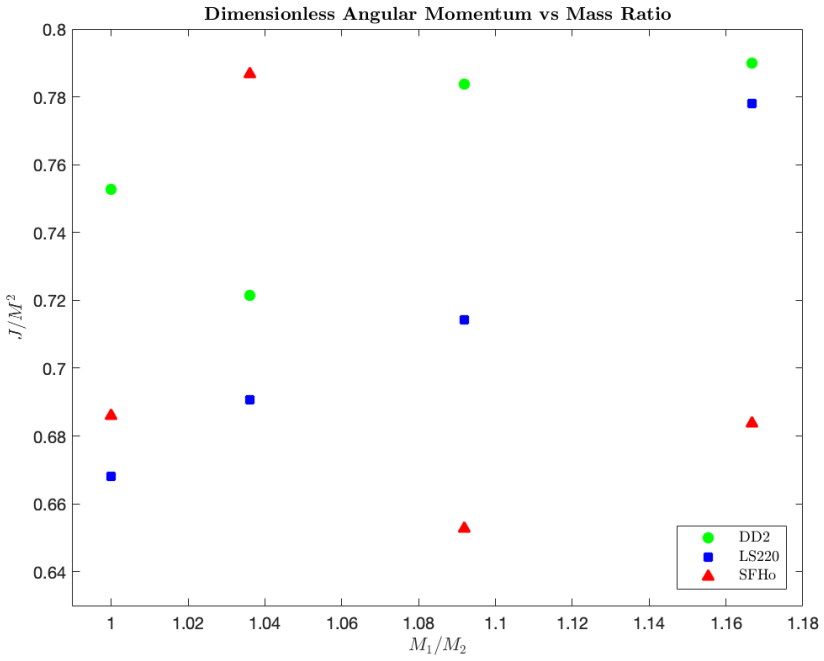

Figure 2-10: The dimensionless binary angular momentum by the end of the simulation plotted against the mass ratio of the components.

Chapter 3

Discussion

3.1 Figure Insights

Looking towards Figure 2-7, we see that for the equations of state DD2 and LS220, the normalized emitted energy E_{GW}/M seems to primarily increase with increasing binary mass M. However, while DD2 features a relatively gentle rise in E_{GW}/M with M, LS220 features a much sharper ascent at lower M before leveling off. Curiously, apart from the lowest M point, the behavior of the SFHo points is essentially a horizontally flipped version of the LS220 points' behavior, as E_{GW}/M appears to briefly increase from $M = 2.600$ $M_\odot$ to $M = 2.615$ $M_\odot$ before plummeting at higher M.

In Figure 2-8 the trends are not as obvious as in Figure 2-7. Indeed, the relationship between E_{GW}/M and the component mass ratio M_1/M_2 is evidently not as simple as is the correspondence between E_{GW}/M and M. DD2's points, for instance, do not display a clear correlation between E_{GW}/M and M_1/M_2. We instead see a steady rise in E_{GW}/M from $M_1/M_2 = 1.000$ to $M_1/M_2 = 1.036$ followed by a seemingly exponential decay. The energies of SFHo express qualitatively similar, but inverted, behavior, plummeting from $M_1/M_2 = 1.000$ to $M_1/M_2 = 1.036$ before rising again sharply and maintaining relatively constant between $M_1/M_2 = 1.092$ and $M_1/M_2 = 1.167$. LS220's points behave the most consistently, as they exhibit a clear negative correlation between E_{GW}/M and M_1/M_2.

Figure 2-9 features development analogous to that of Figure 2-7. DD2's dimensionless angular momenta J/M^2 experience a gradual decline with increasing M, somewhat resembling a vertically flipped version of DD2's trend in Figure 2-7. The points of LS220 and SFHo also follow paths resembling vertically flipped versions of their Figure 2-7 counterparts. LS220's J/M^2 points sharply fall with increasing M before leveling off and slightly rising, and SFHo's points experience a short dip before quickly rising at higher M.

Much like how Figure 2-9 resembles a vertically flipped version of Figure 2-7, so Figure 2-10 resembles a vertically flipped version of Figure 2-8.

This inversion between our E_{GW}/M plots and J/M^2 plots is further supported via Figures 2-4 and 2-5. Notice that as E_{GW}/M rises with time in Figure 2-4, J/M^2 falls in Figure 2-5, Likewise, E_{GW}/M and J/M^2 level off on similar time scales. Such behavior is consistent with that of the gravitational wave amplitude A/M in Figure 2-1, as A/M is greatest during the times of greatest change for E_{GW}/M and J/M^2 and is nearly 0 by the time E_{GW}/M and J/M^2 have basically leveled off. Apparently, at least for binaries, just as the energy flux of a gravitational wave is proportional to the wave amplitude squared

(Schutz 1985), so too is the angular momentum flux, at least approximately. Indeed, I plot the time derivatives of energy and angular momentum against amplitude for DD2-144139 in Figures 3-1 and 3-2 respectively. Although noisy, the plots roughy resemble parabolas, as demonstrated by the superimposed parabolic fits. In this manner, these figures support the hypothesized quadratic relationship between wave amplitude and angular momentum flux.

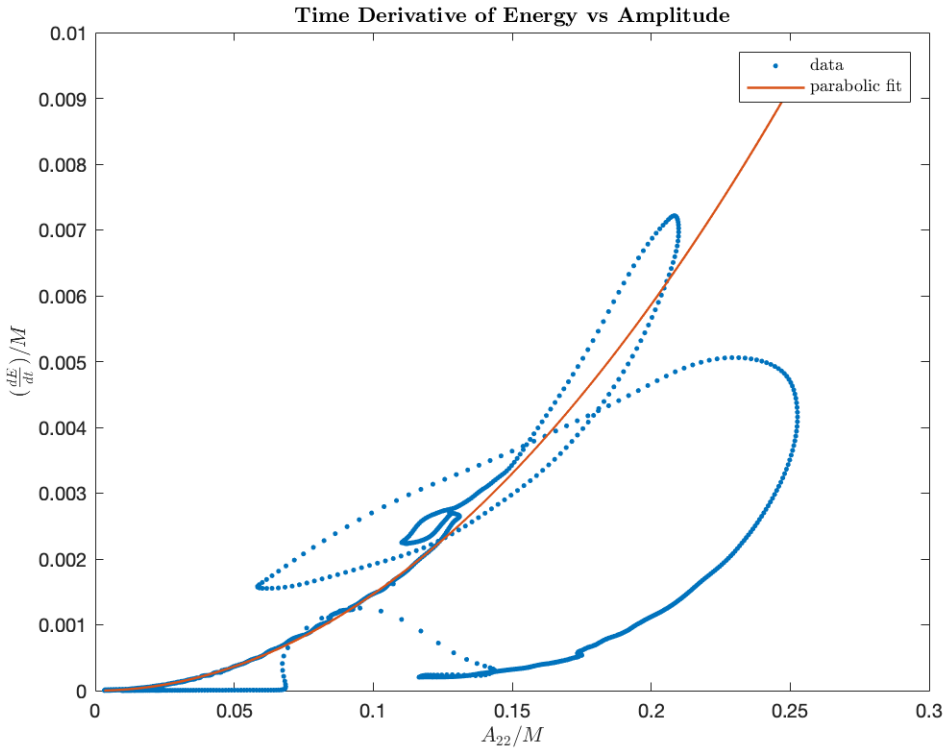

Figure 3-1: The time derivative of the total normalized energy emitted in DD2-144139 via gravitational waves up to time t plotted against the amplitude of the $l = 2$, $m = 2$ mode at time t. The plot bears a roughly parabolic shape and is relatively well fit by the superimposed function $(\frac{dE}{dt})/M = 0.147(A_{22}/M)^2$.

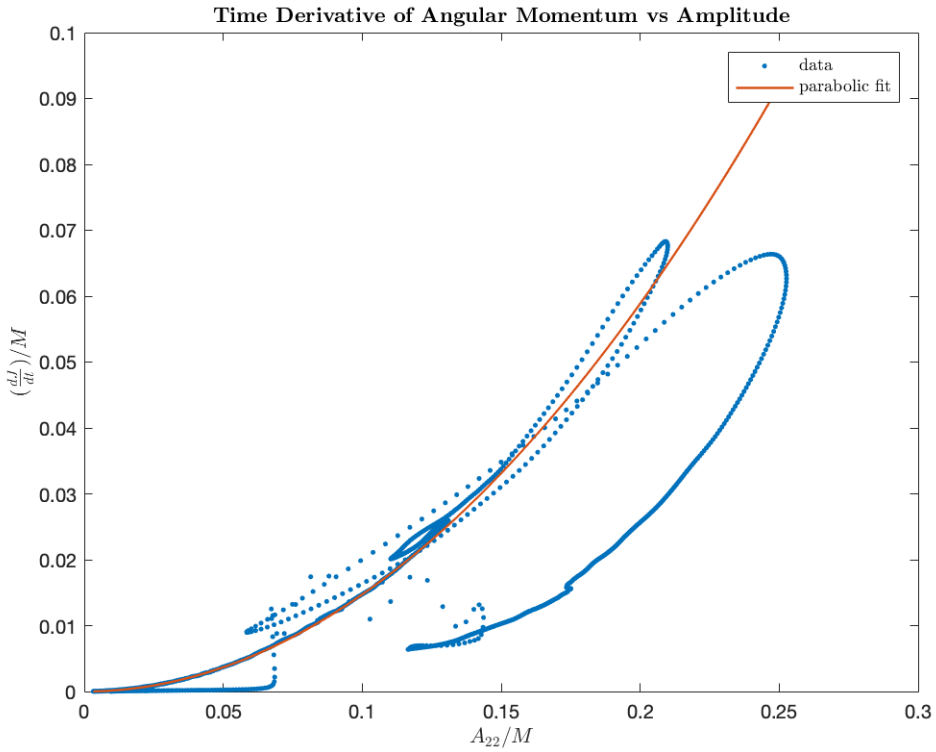

Figure 3-2: The time derivative of the total dimensionless angular momentum emitted in DD2-144139 via gravitational waves up to time t plotted against the amplitude of the $l = 2$, $m = 2$ mode at time t. The plot bears a roughly parabolic shape and is relatively well fit by the superimposed function $(\frac{dJ}{dt})/M^2 = 1.47(A_{22}/M)^2$.

3.2 Final Thoughts

Looking back towards our plots of emitted energy and angular momentum versus binary mass and component mass ratio (Figures 2-7, 2-8, 2-9, and 2-10), we note that any relationships we have inferred are rather robust, as, per our supplemental figures (see Section 4.2) the standard resolution data points tend to fall relatively close to their high resolution counterparts. Furthermore, as covered in our above discussion, we summarize that each equation of state produces significantly distinct relationships between the variables in each figure. This robustness, coupled with the pronounced distinction in behavior between the equations of state, suggests that the relationships demonstrated in these figures could be used to help probe the actual neutron star equation of state. Just as Bernuzzi et al. (2016) asserts that gravitational wave energy measurements from neutron star mergers could be used in conjunction with the relationship displayed in their Figure 3 to place constraints on the equation of state, so too could energy and angular momentum measurements be interpreted via the relationships exhibited in our Figures 2-7, 2-8, 2-9, and 2-10 to accomplish the same goal. Further simulations are needed in order to solidify these relationships. However, just through the work we have done thus far, we have demonstrated yet another way in which WhiskyTHC can further our fundamental understanding of neutron stars.

Chapter 4

Appendix

4.1 TACC Email Correspondence

After months of trying to run the WhiskyTHC software on the Trinity cluster, with little to no support at Trinity, we reached out to a colleague and General Relativity expert, Richard Matzner, at the University of Texas, Austin, about running the code on the Texas Advanced Computing Center (TACC). Prof. Matzner put us in touch with someone familiar with trying to get the software running. The reply is shown here to illustrate the severity of the problem with getting the code to run on any cluster, as well as the seeming futility of using TACC to run our own simulations.

```
From: xxxx
Subject: Re: WhiskyTHC on TACC?
Date: March 13, 2020 at 2:00:35 AM CDT
To: dpooley@trinity.edu
Cc: richard matzner

Prof Pooley,
My apologies on the delay on the reply back here,
but I have been trying to Debug and resolve what exactly is causing the complete and Total Crash
failures on both Stampede2 and Frontera from compiling and running the Verification tests
of the latest "Stable" release from Oct 10, 2019 of the Einstein-Toolkit/Cactus code.

This looks like this "Stable" release from Oct 10 was never properly tested and verified
on Stampede2, so there is No chance we will be able to get this version to work on
Frontera either unless we patch or fix the Errors in this version of the Cactus code and
get it to work on Stampede2 first before trying again to get it to work on Frontera.
Unfortunately Frontera is down right now for Maintenance and Terascale Performance tests
until next week. So we can only Debug these problems right now on Stampede2.

This has been the Major problem we have had to deal with the whole time
on the systems at TACC trying to get the WhiskyTHC software to run ...
and it is Not just problems that you have had with your Linux Cluster.
```

We have to Fight a Constant Battle up at TACC with Very serious Defects and Errors
in the various Release versions of ET/Cactus code which make it almost impossible for us to
get the WhiskyTHC code compiled and built on Top of Cactus to:

1. Run without Crashing with serious Seg Fault, Memory Leak, or other serious Crash Failures.

2. Run Fast and Efficiently on TACC systems so we Do Not waste extremely valuable
    and limited SU Node/Hour compute time resources.

So we have been able to get the WhiskyTHC code to Compile and Run Okay
with a few Release versions of Cactus ...
But only after doing a Lot of Major Fixes and Patches to the Cactus code.

So we can get past condition 1. above, and run the WhiskyTHC without Crashing with Major Failures
on Stampede2 or Lonestar ... but Only After we have done Extensive Bug Fixes and Patches to the Cactus code

But then we have Never been able to pass condition 2 above and be able to
run the WhiskyTHC code in a Very Fast and Efficient High-performance manner on
any of the TACC systems.

This is because there are just Fundamental Defects and Flaws in the Cactus code,
mostly focused in the Carpet AMR Adaptive Mesh Refinement Grid Data management code
inside of the Cactus software framework ...
where the Cactus Software Upgrades to Add-In OpenMP Mutli-threaded code was Never Done properly,
was Never Upgraded to make all the Grid Data management variables 64-bit Safe
(in the Entire History of the Cactus code ... they have Never bothered to
 upgrade the Integer Pointer variables they use to manage Distribute Grid Data storage
 from 32 bit Integer pointers to 64 bit Integer pointer variables until last Spring 2019,
 so there has Always been a 2\^{}32 memory address limit on the number of
 Grid Data storage elements)... and even more Terrible Bad Programming...
They have Never made their code``Thread Safe" so that Multiple Threads in OpenMP
will Not conflict with each other in Shared Data areas in the Grid Data.

So the ET/Cactus developers have been able to just manage to get by over the last Few years
by making lots of ``Quick Fixes" every time people complain about these problems...
on a number of the Older systems that are Important to the Core Developer group at NCSA...

In particular the BlueWaters Cray system at NCSA ...

But they have just Never made any serious efforts to
support the much newer Intel Chip Architectures on the
Lonestar and Stampede2 systems at TACC...
and just Forget it on Frontera ...

They have made Absolutely No Efforts so Far
to try and get Anything with ET/Cactus working on Frontera.

So David Radice has had a Long Battle with the ET/Cactus developers at NCSA
to try and get as many of the Really Bad Defects and Bugs fixed
in the Cactus code ...

So that his WhiskyTHC code will at least Not Crash
and work at some level on the BlueWaters system at NCSA,
and the Comet system at San Diego, and some other
Older systems at other Compute Centers.

But as for getting the WhiskyTHC code to work at TACC...

David spent a Lot of Time Battling with the ET/Cactus developer group at NCSA
to try and get Fixed Many Many Defects and Errors in the Cactus code
that was Just Completely Crashing Everything on the
original Stampede system, and then also
on the Upgraded Stampede2 system with the original KNL nodes.

So that he was able to Finally get the WhiskyTHC code to at least
run in a ``Stable No Crash mode" on the original KNL nodes on Stampede2
using the Cactus 2016 Release to compile and build the WhiskyTHC code...
and he got the 2016 version of Cactus to Finally work
on the KNL nodes so that it Does Not Crash.

and then he Just Stopped and Froze the WhiskyTHC code at that
version of Cactus ... which is at least 4 years old now...
because we think...

He just did Not want to Deal with this Anymore...
The Constant Level of New problems
that get Introduced in Every New Release of the ET/Cactus code.

So this is why in All of David's New Updates to the WhiskyTHC code...
he has Always kept the Compile/Build instructions to Always Stay
with the ET/Cactus 2016 Release.

But Unfortunately, the ET/Cactus 2016 Release will Just Not Work at All
on the newer Stampede2 SKX nodes !!

The new Intel Xeon SkyLake Chip Architecture on the Stampede2 SKX nodes
is just Too New a Chip Architecture ...
and the Cactus 2016 release just Does Not work properly on the SKX nodes,
the Cactus 2016 code just Does Not Auto-Detect properly anything
in the Xeon SkyLake Chip Architecture ...
and the Compute Performance is just Terrible ...
and the WhiskyTHC+Cactus 2016 compiled code just

Crashes All the Time with Memory Leaks and other problems.

So we have had to try and get the WhiskyTHC code to compile
and run with more Recent versions of ET/Cactus ...
where supposedly the support for the New types of Intel Chips is covered in the Cactus code...
But this was also ...
an Extremely Painful amount of Debugging ...
Because All sorts of New Bugs and Defects get Introduced with
Every New Release of the ET/Cactus code.

So the best we have been able to do on the Stampede2 SKX nodes
is to get the WhiskyTHC code to compile with the March 2019 Release of the ET/Cactus code...

and with at least this compiled version of WhiskyTHC + Cactus code ...
We were able to get it to run on the SKX nodes without having to take
Major Crashes all the time ...

But this was only after I had to do some Major Bug Fix patches and Re-writes of the
Carpet AMR code to Solve some Very Serious Memory Leak Errors
that was Crashing the code all the time.

But the Compute Speed performance of this Bug Fix Patched version
of Cactus Spring 2019 code + WhiskyTHC ...

Is Still Totally Terrible !!!

The OpenMP Multi-threading code in the Carpet AMR code that is
Managing the Computations in the Adaptive Data Grid is almost
Completely Non-Functional ...

and it is an Open Question with this version of Cactus ...

If we are Better Off with just Completely Disabling
OpenMP Multi-Threading ... and just use only the
Older MPI Multi-Processing code to manage the
Computations in the Adaptive Data Grid ...

which is then just a Complete Failure to make use
of All the Multi-Core CPU's on these Brand New
Intel Chip Architectures to get a Really Big
Compute Speed Performance Boost from
Adding Open-MP Multi-Threading on top of
MPI Multi-Processing.

So this is about a far a we have been able to get on
the Stampede2 SKX nodes ...

We can make the WhisklyTHC code run at the level
of condition 1. above ...
We can run simulations that Don't Crash for at least 24 hours per job,
that's about how long we can run for each job in a simulation
sequence of chained jobs... so we have to Break-up our
Simulation runs into sequences of many jobs with Checkpoint files
to Stage between jobs ...

So we have been able to get the WhiskyTHC code to run
with a Highly Patched and Bug Fixed Spring 2019 version of Cactus ...
and this Does Not Crash for at Least 24 hours per Job...

But the Compute Speed Performance is Still Absolutely Terrible !!!

We end up Burning Up Lots of Very Valuable Compute Time SU's (Node-Hours)
on Large Simulation jobs with this version of WhiskyTHC+Cactus code...

and this looks to me like...
Just a Terrible Waste of Valuable SU Node-Compute time
to have to be Stuck with this God Awful Slow simulation code.

The OpenMP Multi-threading code in Cactus just Does Not work Properly
on the SKX nodes using this Spring 2019 version of Cactus

So we are just Essentially running Lots of MPI processes on each Node
to use as Many CPU-cores as we can ...

But this means we are Wasting a Lot of CPU cores ...
Not Fully or Properly Using All the CPU cores on each SKX node.

So there is NO Full Utilization of All the CPU cores on each SKX node
running WhiskyTHC with this Spring 2019 version of Cactus.

So this is the best we can do so far with the
WhiskyTHC code on the Stampede2 SKX nodes...
After doing a Whole Lot of work making Bug Fixes to the Cactus Carpet source code.

So I won't even try and Describe here what is going wrong now
with the last ``Stable" Release of Cactus that they put out last Oct 10...

because it looks like they just Rushed this Release Out
without Ever Properly Testing it on the Stampede2 system at TACC.

So we have been Debating this issue back and forth for some time now...

Do we even bother to Start Posting Major Complaints about this
to the Main ET/Cactus Developers Mailing List ...

and Risk getting a Huge of Amount of Critical comments thrown back at us ...
From the Illinois NCSA User community ... about
Do we know what we are doing ...
Or why are we Making Trouble , etc, etc ...

Or do we just give up on this Oct 10 2019 Release...
and just Wait for the Next ``Stable" version
of ET/Cactus to get Released...
and see How Many Bugs the Next version has.

So my Apologies for All this Long List of All the
Total Failures we have found in the Cactus code ...
and All the Hacking of the Cactus source coce
we have to do ...
Just so we can get the WhiskyTHC code to compile and Run
without Crashing on TACC systems.

But I wanted to let you know ...

Your group is Not the only people having trouble
trying to get WhiskyTHC code to compile and Run Okay with ET/Cactus code.

We have been having to Battle with the Cactus code ...
ever since we started working with the WhiskyTHC code.

So the most important thing I wanted to say here is...

We need to Help your group Fix what ever problems
you are having getting the WhiskyTHC code to compile
and Run Okay on your Linux Cluster...

as well as Helping you out with showing what
we have been able to Figure Out on how
to Bug Fix the Damn Cactus code to get
WhiskyTHC so it can work on TACC systems.

The reason why I am saying this is So Important is...

You Need to get the WhiskyTHC software working on
your Linux Cluster as well ...

Because you Need to have A Fall-Back position
where you Can Run some Version of

WhiskyTHC + Cactus on your own Linux Cluster...

and it Runs without Crashing ...

and you Can Run this as Many times as you need to on your own system...
without using up Very Limited and Hard to get SU's up at TACC...

Or Having to Deal with Ridiculous Wait Times of days to weeks on
the Job Queues on Lonestar or Stampede2.

The Compute Resources on both Lonestar and Stampede2 at TACC
have become Completely Over-Committed and
Completely Over-Loaded now ...

So that Any Significant Simulation job sequence I try and Submit
to the ``Normal" Job queue for the SKX Nodes on Stampede2 ...
for even just 20 Node jobs that run for 24 hours each ...

I have to Wait on Average about 7 days in the queue
for each single Job to get started...

There are Typically around 3 to 4 thousand jobs now
just Stacked up on the SKX job queues just Waiting
to run ...
Even on a Good Day !!

So this is A Huge amount of Dead Time just Waiting for each
of your SImulation Jobs to Start...
and this is just the Problem of Wasting Time Doing Nothing
in the Job queues.

The other serious problem is that...

We can Only Run the WhiskyTHC code in a
Very Slow, Very Poor Compute Speed mode...

So we end up Wasting Lots of Valuable Compute Time SU's (Node Hours)...

and trying to get Approval to Renew our Allocations for More SU's
is getting Harder and Harder with NSF now Taking over Everything
on the Review Process for Project Allocations.

So again...
I wanted to List All this ``Problems" with Running things at TACC ...

To make it Clear How Important it is ...

To some how Help your group to get the WhiskyTHC code
to start working on your own Linux Cluster.

You really Need to have the Independence of
Running your simulations ...
at least for small and medium scale Simulation Tests
on your Own Linux Cluster ...

So you are Not getting Dragged down
by All the current Problems at TACC with
a Lot of the Compute Resources on
Stampede2 and Lonestar
are now ...
Completely Overcommitted and Overloaded.

So I have already talked a number of times to the
Senior TACC Admin Staff about How Bad this is
getting on both Lonestar and Stampede2 ...

That the Wait times on the Queues are just getting
Completely Ridiculous !!!

And they have been telling me that there
is No Way they can Solve this problem now...

That NSF has been Telling A Lot of XSEDE Research Project PI's
to Move All their Computing work from NCSA
to the systems at TACC ...
and all the Projects on the Comet system at San Diego SuperComputer Center
to move over to the systems at TACC...
because NSF does Not want to put anymore money into
supporting the BlueWaters Cray system at NCSA
or the Comet System at San Diego, etc, etc....

So TACC now is getting Completely Overloaded now
with Lots of XSEDE Projects moving over from
NCSA and San Diego ... and other XSEDE SuperComputer systems...

And this is All being driven by Decisions made by NSF now
on Funding for Resources on XSEDE Projects.

So anyway ...
TACC has been Advising us to get
Moved Over to Frontera ``As Soon As Possible"
before Frontera gets completely Overloaded Also...

But right now...
the Latest version of Cactus Code is Completely Non-Functional
on either Stampede2 or Frontera.

So I am coming Back Again to this Key point....

We need to Help you also
to get the WhiskyTHC code working on your Linux Cluster !!

You Really Need to have this as a Backup Option
so that you Can Get Some Amount of
Simulation work done with WhiskyTHC on your own
Linux Cluster ...
and Not have to be Totally Dependent on TACC to get work done.

Unfortunately ... We are Stuck in that position right now...
Because the CNS College here had made the decision long ago...
To Not allow or Provide for Any Funding for Compute Clusters
in any of the Dept's in the CNS College.

It may have been possible in the Past to get
Grant money from NSF to Build and Maintain a
Linux Cluster here in the Physics Dept for the
Primary Purpose of doing Numerical GR Simulation work
But Not anymore !!

And as for getting DOE Grant money to do this ...
This is Not Possible Anymore Either !!!

When I was still managing All the Computing Systems in
the Center for Particle Physics here in the Physics Dept...

We had gotten Funding from DOE to Build and Manage
a Resonably Large 250+ Node Linux Cluster system
to Support the BaBar, MINOS, SNO, and RHIC STAR experiment groups here...

But DOE then completely Canceled and Shutdown All the Funding to Maintain this System..
and we had to Tear Down this whole system well over 10 years ago...

So this is Not Possible Anymore in the Physics Dept here at UT Austin
to Build Any Kind of Large New Linux Cluster system here in the Dept
So that we have some Backup Option to
Run SImulations here in the Physics Dept.

But if you on the other hand ...

Do Have some reasonable size Linux Cluster
in your own Group ...

You should Not give up on Trying to get the
WhiskyTHC code to Finally Work on your system...

You then have some Fall-back options
to get SImulation work done on your own Linux Cluster system.

Things up at TACC are just getting way too Overloaded
with Too many Jobs constantly Piling Up on the
Job queues.

So the Final point to get to here is ...

Can you get me more information
on What Exactly are the Types of
Linux Nodes you have in your Cluster...

That is what type of Intel or AMD processors your Nodes have,
and how much Memory is on each Node...
and what type of Linux OS are you using on each Node.

Or can you get me in contact with who ever is your
Sys Admin for your Linux Cluster ...

So I can Figure Out How we can Help you to get
the WhiskyTHC code to Work on your Linux Cluster.

It may turn out that the only issues are just
to Figure out which Version of Cactus, and
which version of GNU compilers, and which version
of the other software packages are needed ...

Just trying to get the right combination of software
packages to Fit together ...
So the WhiskyTHC software will Finally Compile
and Run Okay on your Linux system.

Then you have a Fall-Back solution...
that you can Always Run WhiskyTHC simulations
on your own Linux Cluster nodes...
For what ever size simulation your Cluster can handle...

And then you are Not Stuck on being Totally Dependent
on using only the TACC systems ...

which is what we are stuck with right now.

This is Really A Sore point for me right now...

That we have No Computing Systems here in the Physics Dept
that can be used to get Any Work Done on Simulations.

We are Totally Dependent on TACC to get this kind Simulation work Done.

So I think what ever Money and Time you Invested
in Building your own Linux Cluster ...

This is still worth it to try and get the WhiskyTHC code
to work on your System ...

and I think this is definitely a good idea for us to Help
you out on this if we can,

Otherwise you are stuck in the same position we are in now.

Which is we are Totally Dependent on TACC now
with what are all the problems and limitations
that NSF is Forcing Everybody to Deal with up at TACC now.

So let me know if there is anyone else
I should contact who is your Sys Admin,
or Grad student, or who ever is maintaining
your Linux Cluster ...

I would like to Find out what Exactly
is the Hardware and Software you have
on your Linux Nodes ... and then maybe
we can Help you to get things Fixed with
the WhiskyTHC software ...
So it can actually Run Okay on your Linux Cluster also.

4.2 Supplementary Figures

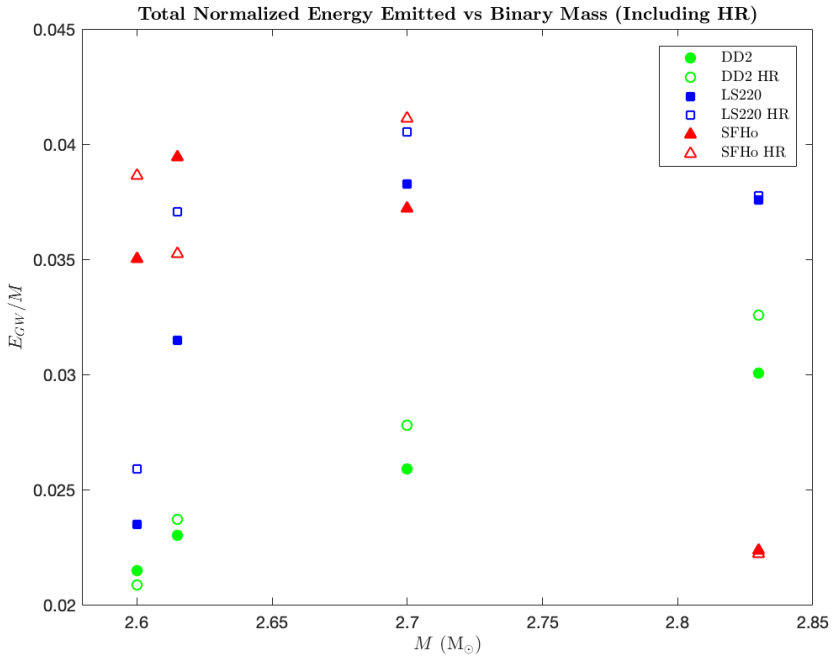

Figure 4-1: The total normalized energy emitted by a binary via gravitational waves by the end of the simulation, including energy radiated before the simulation began, plotted against the binary mass. Data from high resolution simulations are included.

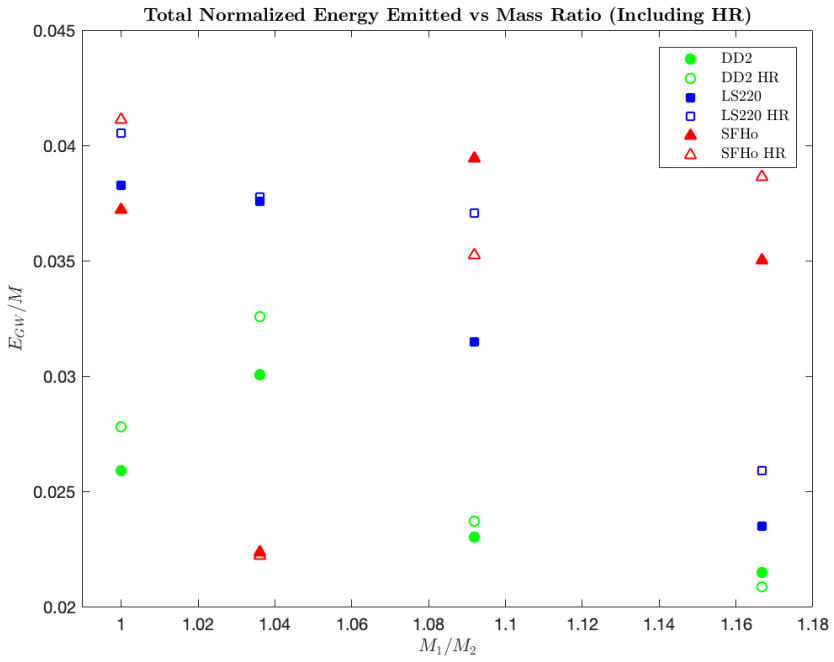

Figure 4-2: The total normalized energy emitted by a binary via gravitational waves by the end of the simulation, including energy radiated before the simulation began, plotted against the mass ratio of the components. Data from high resolution simulations are included.

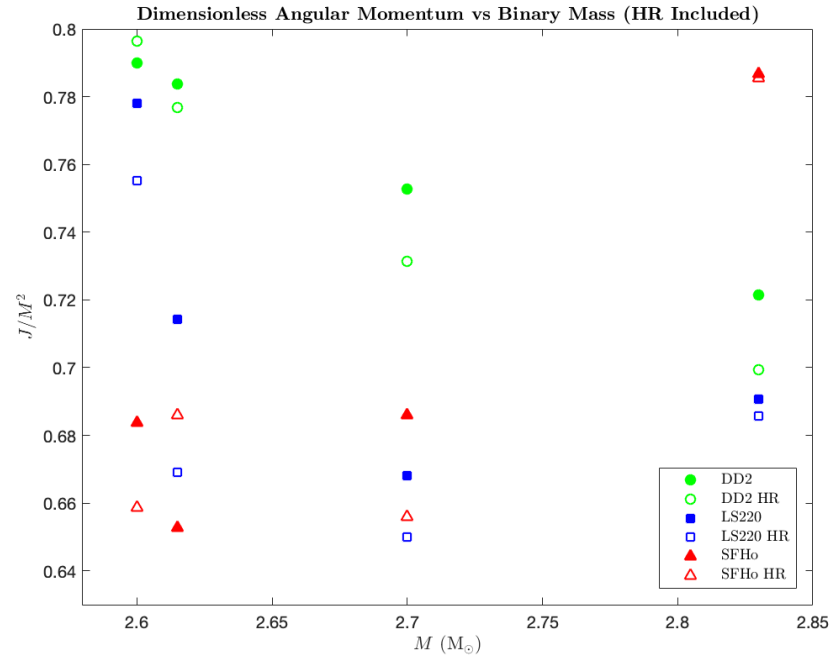

Figure 4-3: The dimensionless binary angular momentum by the end of the simulation plotted against the binary mass. Data from high resolution simulations are included.

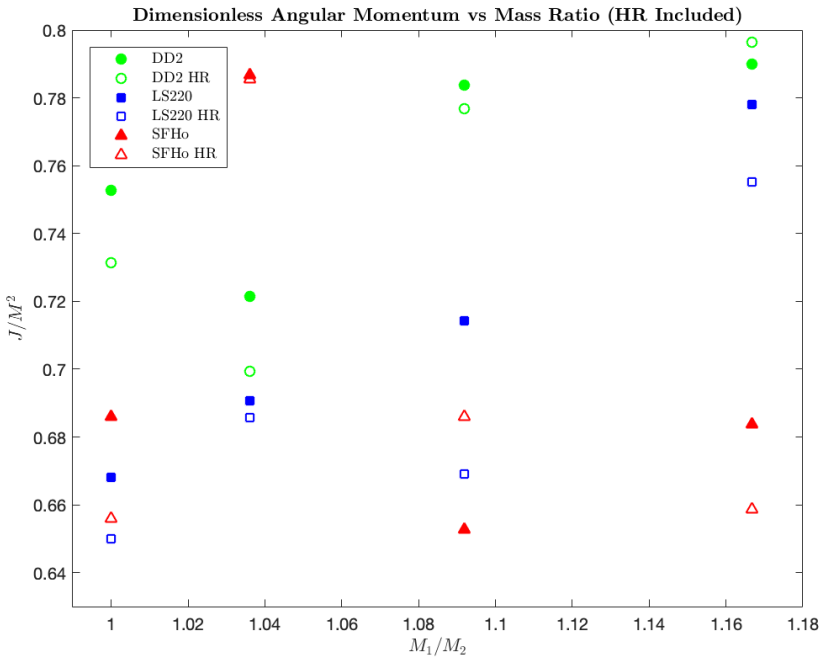

Figure 4-4: The dimensionless binary angular momentum by the end of the simulation plotted against the mass ratio of the components. Data from high resolution simulations are included.

Bibliography

Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016, PRL, 116, 061102

Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017, PRL, 119, 161101

Bernuzzi, S., Dietrich, T., & Nagar, A. 2015, PRL, 115, 091101

Bernuzzi, S., Radice, D., Ott, C. D., et al. 2016, PRD, 94, 024023

Bradt, H. 2008, Astrophysics Processes

Burbidge, E. M., Burbidge, G. R., Fowler, W. A., et al. 1957, Reviews of Modern Physics, 29, 547

Cameron, A. G. W. 1957, PASP, 69, 201

Fryer, C. L., Belczynski, K., Ramirez-Ruiz, E., et al. 2015, ApJ, 812, 24

Lattimer, J. M. 2012, Annual Review of Nuclear and Particle Science, 62, 485

Metzger, B. D. 2017, arXiv e-prints, arXiv:1710.05931

Miller, M. C., & Yunes, N. 2019, Nature, 568, 469

Radice, D., Rezzolla, L., & Galeazzi, F. 2014, Classical and Quantum Gravity, 31, 075012

Raithel, C. A., Özel, F., & Psaltis, D. 2018, ApJ, 857, L23

Schutz, B. F. 2009, A First Course in General Relativity

Taylor, E.-F., Wheeler, J.-A., & Bertschinger, E. 2019