Observable Signatures of General Relativistic Dynamics in Compact Binaries

by

Ryan Nathan Lang

S.B. Electrical Science and Engineering
Massachusetts Institute of Technology, 2002

S.B. Physics
Massachusetts Institute of Technology, 2003

M.Eng. Electrical Engineering and Computer Science
Massachusetts Institute of Technology, 2003

Submitted to the Department of Physics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2009

Author .....
Department of Physics
May 22, 2009

Certified by .....
Scott A. Hughes
Associate Professor of Physics
Thesis Supervisor

Accepted by .....
Thomas J. Greytak
Lester Wolfe Professor of Physics
Associate Department Head for Education

Observable Signatures of General Relativistic Dynamics in Compact Binaries

by

Ryan Nathan Lang

Submitted to the Department of Physics
on May 22, 2009, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

The effects of general relativity (GR) in astrophysical systems are often difficult to calculate, but they can have important consequences for observables. This thesis considers the impact of previously-ignored GR effects in two different types of compact binary systems. The first is the coalescence of massive black holes in high-redshift galaxies. The gravitational waves (GWs) from these systems can be detected by the proposed low-frequency gravitational wave detector LISA and used to determine the various parameters which characterize the binary.

Most studies of LISA's parameter estimation capability have ignored a significant piece of physics: the relativistic precession of the binary's angular momentum vectors. In the first two-thirds of this thesis, we show how including precession effects in the waveform model helps to break various degeneracies and improve the expected parameter errors. We give special attention to the localization parameters, sky position and distance. When distance is converted to an approximate redshift, these parameters define a "pixel" on the sky in which astronomers can search for an electromagnetic counterpart to the GW event.

The final third of this thesis focuses on stellar-mass compact binaries in which at least one member is a neutron star. The measurement of tidal effects in these systems may shed some light on the poorly understood high-density equation of state. We first calculate the point at which a neutron star tidally disrupts in the field of a black hole. Previous calculations of this effect have used Newtonian self-gravity, which is inappropriate for a neutron star; we correct this by using relativistic perturbation theory. We then turn to small tidal distortions of neutron stars, which can be characterized by a quantity known as the Love number. We calculate relativistic Love numbers for a wide variety of equations of state and investigate their impact on the GWs from neutron star-neutron star binaries.

Thesis Supervisor: Scott A. Hughes
Title: Associate Professor of Physics

Acknowledgments

This thesis is the culmination of eleven fantastic years at MIT. Even just focusing on the past six, there are many, many people to thank! Before I begin, though, I must first thank God, who has surrounded me with all of these fantastic people, as well as given me the ability and opportunity to do what I want with my life.

This thesis would not exist without the support and guidance of my advisor, Scott Hughes. I feel very lucky to have had Scott as an advisor, since he is an excellent teacher and researcher. In addition, despite his ridiculously busy schedule, he is always available to answer questions or just to chit-chat about some random topic. Most of all, I am grateful to Scott for taking me in as his student when I was pretty miserable in my previous research group. If I had not found Scott when I did, I might not be getting a Ph.D. at all.

I would also like to thank the rest of my committee, Ed Bertschinger and Nergis Mavalvala, for all of their advice and support, especially when the focus of the last thesis chapter changed rather suddenly. Neil Cornish (of Montana State University) was also very helpful in my research, even though our joint projects were not ready in time for this thesis. Hopefully we will soon finish them! Samaya Nissanke is our partner-in-crime on those projects; however, her greatest contribution to my thesis was as a wonderfully supportive friend. I look forward to many years of hanging out at conferences together.

I would be remiss if I did not mention the funding for all of these projects. The work in this thesis was supported by NASA grants NAGW-12906, NNG05G105G, and NNX08AL42G, as well as NASA contract no. 1291617 and NSF grant PHY-0449884. In the last year, I have been a Bruno Rossi Fellow, supported by the generous contributions of Jim and Sylvia Earl. I am really appreciative for this fellowship, since it allowed me to focus on thesis at a critical time without the distraction of teaching.

Coming to work each day was always enjoyable due to the great people, including my academic siblings (Pranesh Sundararajan, Sarah Vigeland, Stephen O'Sullivan, and Leo Stein). I can't wait to plan "Scott-fest" someday down the road! I want to

give a special thanks to my fellow occupants of 37-602, Josh Carter, Molly Swanson, Chris Williams, Leslie Rogers, and Leo. They all made grad school a very fun (if not always productive) time with games of Scrabulous, Traveler IQ Challenge, Indoor Paper Baseball, and Cosmic Office Volleyball, as well as discussions of life, politics, and Lost. I especially want to thank Josh, a fantastic guy whom I am lucky to have as one of my best friends.

Outside of work, I have been blessed to know many fantastic people. Mat Willmott is the best friend that anyone can ask for, and I will miss him (and our Friday lunches) very much next year. Rich Redemske, my freshman year roommate, is one of my oldest and greatest friends, and his wife Alicia is also incredibly awesome (and not just because is she is a Republican). Like Rich, Zhenye Mei has been there since the beginning (New House 2), and he is still always there for me when I need to talk about baseball or something more serious. Victoria Anderson (along with Mr. Cow) was a fantastic roommate during my first year of grad school, and even though she moved away, she still amuses me daily with IM conversations about all the people we can't stand. I did it for THE TEAM one last time, Vic. I don't know how my current roommate, Dan Nezich, managed to put up with me for five years, but we did have a lot of fun watching the antics of Donald Trump and Jeff Probst. Dan Kwon and I have experienced MIT together for the entire eleven years, starting in concert band, moving on to IM softball and football, and finally graduating together (again) in just a few weeks. Alisa Marshall is my fantastic Mystery Hunt co-captain and close confidant in many gym and Windsor lot conversations. Her fiancé David Kovacs is an awesome guy, and not just because he loves all the same TV that I do. Keith Santarelli can be a bit of a jerk, but he did buy me champagne for my defense, so I guess I can deal. (Kidding, Keith! You rock.) Rich Ott gets special thanks for keeping me sane during our first year of grad school, when we were both in the same unhappy research situation. Dennis Wu graduated about a zillion years ago and left for Asia, which was sad because he is a great guy. Cindy Jao and I go way back to Next House days, and luckily we have remained friends since. Along the way, I've managed to befriend her awesome fiancé, Brian Wu. Ian Lai and I go

even further back: We were temp roommates, and eleven years later, he made it to my thesis defense. Lokie Cerritelli has been over my apartment many evenings and is always up for good conversation or maybe just some much-needed venting about work. Joe Wilkins is a hilarious guy, and his new wife Bonnie Andrews might be the nicest person alive. And finally, I have to give a shoutout to the greatest trivia team ever, the Puffywumps, including (relatively) new but awesome friends Cathy Hoelscher, Jeff Simpson, and Andrew Freeman. And of course, the guy who came up with the famous name, the one-and-only Ken Marr. Thanks to all these people for making my time in Boston so happy.

Most of all, though, I want to thank my wonderful family. My parents, Greg and Lisa Lang, are the two most fantastic people on the planet. Their love and support, even when they might have questioned this whole "grad school thing," have been invaluable. They have sacrificed all their lives so that I can have what I need to succeed. My sister Rachel is also an awesome person who understands me like no one else in the world. I have been especially fortunate to have her close by in New York City the last few years. I have also been lucky to have my wonderful grandparents Peter and Ethel Lang. It was Grandpa Lang who encouraged my interest in science as a child, and both of them have always been there for me ever since. My aunts, uncles, and cousins have also been incredibly supportive and loving. Finally, I want to dedicate this thesis to my grandparents LaVerne (Tim) and Betty Utley. Their financial support made it possible to come to MIT, but it was their incredible friendship and love which helped me finish. I wish they could be here to see me graduate, but I still feel them watching over me. I love you, Gam and Gamp.

Contents

1 Introduction **23**

1.1 Einstein's theory of general relativity 25

1.1.1 The Equivalence Principle and its consequences 25

1.1.2 Key features of general relativity 27

1.2 Compact objects and binaries 34

1.2.1 White dwarfs 34

1.2.2 Neutron stars 35

1.2.3 Black holes 36

1.2.4 Compact binaries 39

1.3 Gravitational waves 42

1.3.1 Basic theory of gravitational waves 44

1.3.2 Gravitational wave sources and detection 47

1.4 Outline of this thesis 51

2 Measuring massive black hole binaries with gravitational waves **55**

2.1 Introduction 55

2.1.1 Background to this analysis 55

2.1.2 Black hole spin and spin precession 58

2.1.3 This analysis 61

2.1.4 Outline of this chapter 64

2.2 Gravitational waves from binary black hole inspiral 66

2.2.1 Intrinsic waveform 67

2.2.2 Precession equations 71

2.2.3 Extrinsic effects . . . 73
2.3 Measurement and parameter estimation with LISA . . . 78
2.3.1 Theory . . . 78
2.3.2 LISA detector and astrophysical noise . . . 82
2.4 Results . . . 84
2.4.1 Procedural issues . . . 84
2.4.2 Black hole masses and spins . . . 90
2.4.3 Sky position and distance to source . . . 98
2.5 Summary and conclusions . . . 102

3 Localization of massive black hole binaries 113
3.1 Introduction . . . 113
3.1.1 The LISA GW pixel . . . 113
3.1.2 Electromagnetic counterparts to MBHB GW events . . . 115
3.1.3 Outline of this chapter . . . 117
3.2 The LISA pixel at merger . . . 119
3.3 Time and position dependence of the LISA pixel . . . 123
3.3.1 Summary of KHMF . . . 124
3.3.2 Results I: Time evolution of localization accuracy . . . 128
3.3.3 Results II: Angular dependence of localization accuracy . . . 139
3.4 Summary and conclusions . . . 145

4 Tidal perturbations of neutron stars 151
4.1 Introduction . . . 151
4.1.1 Neutron stars and the equation of state . . . 151
4.1.2 Outline of this chapter . . . 154
4.2 Newtonian models of tidally distorted stars . . . 156
4.3 Metric of a perturbed neutron star . . . 160
4.3.1 Unperturbed neutron star . . . 160
4.3.2 General perturbations to a spherically symmetric spacetime . 162
4.3.3 Even parity perturbation equations . . . 165

4.3.4 Boundary conditions: origin . . . 167
4.3.5 Boundary conditions: large r . . . 168
4.3.6 Boundary conditions: surface . . . 169
4.4 Metric of a neutron star in a binary . . . 172
4.4.1 Tidal field of the black hole . . . 173
4.4.2 Matching procedure . . . 176
4.5 Results: Tidal disruption . . . 178
4.5.1 Procedural issues . . . 178
4.5.2 Comparison to other relativistic tidal problems . . . 181
4.5.3 Comparison to classical tidal problems . . . 184
4.6 Love numbers . . . 189
4.6.1 Effect of tidal excitations on GW phase . . . 190
4.6.2 Calculating the Love number . . . 193
4.7 Results: Love numbers . . . 195
4.7.1 Polytropes . . . 195
4.7.2 General equations of state: definitions . . . 197
4.7.3 General equations of state: Love numbers . . . 207
4.8 Summary and conclusions . . . 214

List of Figures

2-1 These figures depict the "polarization amplitude" $A_{\rm pol}(t)$ of a signal measured in detector I for a selection of spins: $\chi_1 = \chi_2 = 0$ (blue line), $\chi_1 = \chi_2 = 0.1$ (red line), $\chi_1 = \chi_2 = 0.5$ (black line), and $\chi_1 = \chi_2 = 0.9$ (green line). ($\chi = |\mathbf{S}|/m^2$ is the dimensionless spin parameter.) The top figure shows the final two years of inspiral. The spinless curve has periodicity of one year, corresponding to the motion of LISA around the Sun. Notice that as spin is introduced, the curves become more strongly modulated, with the number of additional oscillations growing as the spin is increased. By tracking these spin-precession-induced modulations, it becomes possible to better measure parameters like mass and sky position and measure spin for the first time. The bottom figure shows a close-up of the final months of inspiral. Precession effects increase drastically as merger approaches. 59

2-2 Distribution of errors in chirp mass $\mathcal{M}$ for 10^4 binaries with $m_1 = 10^6 M_\odot$ and $m_2 = 3{\times}10^5 M_\odot$ at $z = 1$. The dashed line is the precession-free calculation; the solid line includes precession. Precession reduces the measurement error by about an order of magnitude. 90

2-3 Distribution of errors in reduced mass μ for 10^4 binaries with $m_1 = 10^6 M_\odot$ and $m_2 = 3{\times}10^5 M_\odot$ at $z = 1$. The dashed line is the precession-free calculation; the solid line includes precession. Precession has an enormous effect on the reduced mass, which was previously highly correlated with the parameters β and σ. 91

2-4 Distribution of errors in individual hole masses for 10^4 binaries at $z = 1$. The solid line is $m_1 = 10^6 M_\odot$, while the dashed line is $m_2 = 3\times10^5 M_\odot$. The individual masses are not determined as well as $\mathcal{M}$ and μ, but they are better behaved parameters when precession is introduced. 92

2-5 Distribution of errors in dimensionless spin parameters χ_1 (solid line) and χ_2 (dashed line) for 10^4 binaries with $m_1 = 10^6 M_\odot$ and $m_2 = 3 \times 10^5 M_\odot$ at $z = 1$. In each binary, the spin values are randomly selected between 0 and 1. The higher mass then has, on average, higher total spin and more effect on the precession. 93

2-6 Distribution of errors in dimensionless spin parameter χ_1 for 10^4 binaries with $m_1 = 10^6 M_\odot$ and $m_2 = 3 \times 10^5 M_\odot$ at $z = 1$. Here, spin magnitudes have been set to a specified value — low spin, $\chi_1 = \chi_2 = 0.1$ (dashed line), and high spin, $\chi_1 = \chi_2 = 0.9$ (solid line). Since greater spin more strongly impacts the waveform, the high spin case is measured more accurately. 94

2-7 Distribution of the major axis of the sky position error ellipse, $2a$, for 10^4 binaries with $m_1 = 10^6 M_\odot$ and $m_2 = 3 \times 10^5 M_\odot$ at $z = 1$. The dashed line is the precession-free calculation; the solid line includes precession. Sky position, as an extrinsic parameter, is improved somewhat indirectly by precession; therefore, the improvement is less than for the masses. 100

2-8 Distribution of the minor axis of the sky position error ellipse, $2b$, for 10^4 binaries with $m_1 = 10^6 M_\odot$ and $m_2 = 3 \times 10^5 M_\odot$ at $z = 1$. The dashed line is the precession-free calculation; the solid line includes precession. 101

2-9 Distribution of errors in the luminosity distance for 10^4 binaries with $m_1 = 10^6 M_\odot$ and $m_2 = 3 \times 10^5 M_\odot$ at $z = 1$. The dashed line is the precession-free calculation; the solid line includes precession. 101

3-1 Distribution of the major axis $2a$ (top) and minor axis $2b$ (bottom) of the sky position error ellipse for 10^4 binaries with $m_1 = 10^6 M_\odot$ and $m_2 = 10^5 M_\odot$ (solid line), $3 \times 10^5 M_\odot$ (dashed line), and $10^6 M_\odot$ (dash-dotted line) at $z = 1$. Note the bimodal character of the major axis as well as the long tail down to small minor axis, both of which are particularly prominent for larger mass ratios. 122

3-2 Evolution of the sky position error ellipse for nine individual binaries selected from a set of 10^4. All have $m_1 = 10^6 M_\odot$, $m_2 = 3 \times 10^5 M_\odot$, and $z = 1$. The ellipses are oriented so their major axes are parallel to the x-axis and their minor axes are parallel to the y-axis; the axes are labeled in arcminutes. From outside in, the ellipses are evaluated at 28, 21, 14, 7, 4, 2, 1, and 0 days before merger. 130

3-3 Same as Fig. 3-1, but with $m_1 = 10^6 M_\odot$, $m_2 = 3 \times 10^5 M_\odot$, and $z = 1$ at different values of N (the number of days before merger). Reading from left to right, $N = 0$ (solid line), 1 (dashed line), 7 (dotted line), and 28 (dash-dotted line). Clearly, the largest change — in shape as well as median — happens between merger and one day before. . . . 132

3-4 Medians of the sky position ellipse axes for Monte Carlo runs of 10^4 binaries as a function of time before merger. Major axes $2a$ are on the left; minor axes $2b$ are on the right. Data were only output at the marked points; the lines are there just to guide the eye. The masses have been subdivided into "low," "intermediate," and "high" groups; the exact values (in units of solar masses) are given in the legends. Note also the different scales (arcminutes and degrees) on the y-axis. 134

3-5 Medians of $2a$ and $2b$ as a function of time, comparing an analysis that accounts for spin-induced precession to one that neglects it. Solid lines trace the evolution of $2a$; dashed lines trace $2b$. Precession results are marked with crosses, no precession with circles. The left plot shows a low-mass case, $m_1 = 3 \times 10^5 M_\odot$ and $m_2 = 10^5 M_\odot$; the right plot shows an intermediate-mass case, $m_1 = m_2 = 10^6 M_\odot$. Both plots are for $z = 1$. 136

3-6 Off-diagonal covariance matrix entries illustrating correlation between sky position and binary orientation as a function of time, for the binary in Fig. 3-2a. The correlations decrease rapidly in the final day before merger, when precession effects are maximal. 137

3-7 Distribution of $\Delta D_L/D_L$ for 10^4 binaries with $m_1 = 10^6 M_\odot$, $m_2 = 3 \times 10^5 M_\odot$, and $z = 1$ at different values of N (the number of days before merger). Reading from left to right, $N = 0$ (solid line), 1 (dashed line), 7 (dotted line), and 28 (dash-dotted line). 138

3-8 Dependence of the localization errors on $\bar{\mu}_N$. The major axis $2a$ of the sky position error ellipse is on the top left, the minor axis $2b$ on the top right, and the luminosity distance errors $\Delta D_L/D_L$ on the bottom. Each datum represents the median of 10^4 binaries with $m_1 = 10^6 M_\odot$, $m_2 = 3 \times 10^5 M_\odot$, and $z = 1$; all other parameters are selected randomly (except for $\bar{\mu}_N$, whose range is limited to the bin width). 142

3-9 Dependence of the localization errors on $\bar{\phi}_N$. The major axis $2a$ of the sky position error ellipse is on the top left, the minor axis $2b$ on the top right, and the luminosity distance errors $\Delta D_L/D_L$ on the bottom. Each datum represents the median of 10^5 binaries with $m_1 = 10^6 M_\odot$, $m_2 = 3 \times 10^5 M_\odot$, and $z = 1$; all other parameters are selected randomly (except for $\bar{\phi}_N$, whose range is limited to the bin width). 144

3-10 Sky maps of major axis $2a$ (top row, in arcminutes), minor axis $2b$ (second row, in arcminutes), localization ellipse area $\Delta\Omega_N$ (third row, in square degrees) and $\Delta D_L/D_L$ (bottom row) for LISA observations of MBH binaries in different parts of the sky. Data in the left column are presented in ecliptic coordinates; data on the right are in Galactic coordinates, with the Galactic center at the middle. Note that the level of $\bar{\phi}_N$ variation is so small that it would not show up in these figures; accordingly, they are essentially just remappings of Fig. 3-8. 146

4-1 Comparison of perturbative GR results and the classical Jeans results (modified for a compressible star) for the tidal distortion and disruption of a 0.49 $M_\odot$ white dwarf orbiting a 3 $M_\odot$ black hole. 186

4-2 Comparison of perturbative GR results and the classical Roche results (modified for a compressible star) for the tidal distortion and disruption of a 0.49 $M_\odot$ white dwarf orbiting a black hole. The top plot shows the case $m_{\rm BH} = 100\ M_\odot$ ($q \to \infty$), while the bottom plot shows the case $m_{\rm BH} = 0.49\ M_\odot$ ($q = 1$) 187

4-3 Love number k_2 as a function of compactness $m_{\rm NS}/R$. The solid line depicts k_2 for the equation of state $P = 7.4 \times 10^4 \rho_0^2$. The points marked $\times$ are Hinderer's original, incorrect results for a $\gamma = 2$ ($n = 1$) polytrope [116]. The dashed line is the $n = 1$ result from Berti, Iyer, and Will [31], who use a rotating star to approximate a tidally deformed one. 196

4-4 Complete equation of state for cases with normal matter at the surface. Starting from low densities, the various pieces are: FMT (dotted line), BPS (dashed line), BBP (thick solid line), NV (dot-dashed line), and various high density EOS (thin solid lines). The following figures provide close-ups on these high density equations of state, as well as those including quark matter which do not have this low-density behavior. . 203

4-5 Top: High density equations of state provided by Jim Lattimer and which include only nucleons. Bottom: Mass-radius relations for these equations of state. 204

4-6 Top: High density equations of state provided by Jim Lattimer and which include more exotic components, including pions (PS), hyperons (GM1-GM3), kaons (GS1-GS2), hyperons and free quarks (PCL2), and strange quark matter (SQM1-SQM3). While most of these equations of state match to the low density EOS pictured in Fig. 4-4, the SQM EOS reach zero pressure at a large finite density. Bottom: Mass-radius relations for these equations of state. 205

4-7 Top: High density equations of state provided by Mark Alford. APR includes only nucleons (except for a pion condensation phase transition), while all the others include quark matter as described in the text. All but ALF3C match to the low density EOS pictured in Fig. 4-4. Bottom: Mass-radius relations for these equations of state. . . . 206
4-8 Love number k_2 as a function of compactness $m_{\rm NS}/R$ for Lattimer equations of state which include only nucleons. 208
4-9 Love number k_2 as a function of compactness $m_{\rm NS}/R$ for Lattimer equations of state which include exotic components (excluding SQM1-SQM3). 209
4-10 Love number k_2 as a function of compactness $m_{\rm NS}/R$ for equations of state provided by Mark Alford (excluding ALF3C). 209
4-11 Love number k_2 as a function of compactness $m_{\rm NS}/R$ for equations of state SQM1-SQM3 and ALF3C. 210
4-12 Number of GW cycles $\Delta N_{\rm cyc}$ contributed by tidal term for Lattimer equations of state which include only nucleons. Each neutron star in the binary has mass $m_{\rm NS}$. 212
4-13 Number of GW cycles $\Delta N_{\rm cyc}$ contributed by tidal term for Lattimer equations of state which include exotic components. Each neutron star in the binary has mass $m_{\rm NS}$. 213
4-14 Number of GW cycles $\Delta N_{\rm cyc}$ contributed by tidal term for equations of state provided by Mark Alford. Each neutron star in the binary has mass $m_{\rm NS}$. 213

List of Tables

2.1 Median errors in intrinsic quantities for 10^4 binaries of various masses at $z = 1$, including comparisons with the "no precession" case where possible. We have omitted the errors in chirp mass and reduced mass for equal-mass binaries because that parameterization of the waveform fails the Gaussian approximation at those points. 95

2.2 Median errors in intrinsic quantities for 10^4 binaries of various masses at $z = 3$. 96

2.3 Median errors in intrinsic quantities for 10^4 binaries of various masses at $z = 5$. The results for the highest masses are meaningless — the parameters are completely undetermined. 97

2.4 Median errors in sky position for 10^4 binaries of various masses at $z = 1$, including comparisons with the "no precession" case. Note that the given major axis, minor axis, and ellipse area are the medians for each data set and do not correspond to the same binary. However, they still represent an average sky position error ellipse in the following sense: $\sqrt{\pi ab}$, calculated using the median values of $2a$ and $2b$, differs in most cases by less than 10% from the median value of $\sqrt{\Delta\Omega_N}$ (except at more extreme mass ratios — when $m_1/m_2 = 10$, the difference can be 25%). 103

2.5 Median errors in luminosity distance for 10^4 binaries of various masses at $z = 1$, including comparisons with the "no precession" case. 104

2.6 Median errors in sky position for 10^4 binaries of various masses at $z = 3$. 105

2.7 Median errors in luminosity distance for 10^4 binaries of various masses at $z = 3$. 106

2.8 Median errors in sky position for 10^4 binaries of various masses at $z = 5$. Again, the results for the highest masses are essentially meaningless—the parameters are completely undetermined. 107

2.9 Median errors in luminosity distance for 10^4 binaries of various masses at $z = 5$. 108

3.1 Example of errors (diagonal elements) and correlations (off-diagonal elements) for a binary with $m_1 = 3 \times 10^6\, M_\odot$ and $m_2 = 10^6\, M_\odot$ at $z = 1$. The errors in $\bar{\phi}_L$, $\bar{\phi}_N$, and Φ_c are measured in radians; the error in t_c is measured in seconds. This example was taken from the same Monte Carlo distribution used to make Table 3.3; in this particular case, the randomly distributed parameters have the values $\cos\bar{\theta}_L = -0.628$, $\cos\bar{\theta}_N = 0.850$, $\bar{\phi}_L = 3.50$ rad, $\bar{\phi}_N = 0.514$ rad, $t_c = 6.90 \times 10^7$ s, $\beta = 1.48$, and $\sigma = 0.107$. Entries containing $\cdot$ can be found by symmetry. 126

3.2 Example of errors and correlations for the same binary as shown in Table 3.1 if correlations between intrinsic and extrinsic parameters are neglected. 127

3.3 Ten Monte Carlo points comparing sky position accuracy calculated using the full Fisher matrix technique to that calculated using the KHMF approximation. All ten binaries have $m_1 = 3 \times 10^6\, M_\odot$ and $m_2 = 10^6\, M_\odot$ at $z = 1$. 128

3.4 Typical ranges of sky position and distance measurement accuracy as a function of time until merger for binaries at $z = 1$. The total mass M has been divided into three categories: low mass ($M \lesssim 10^6 M_\odot$), intermediate mass ($10^6 M_\odot \lesssim M \lesssim 4 \times 10^6 M_\odot$), and high mass ($M \gtrsim 6 \times 10^6 M_\odot$). Angles are in degrees except for those marked "a," which are in arcminutes; solid angles are always in square degrees. For rows marked "b," some very massive systems are excluded from the data. In those cases, the position and distance are very poorly constrained that far in advance of merger. In some cases, the binary is even out of band. 139

3.5 Typical ranges of sky position and distance measurement accuracy as a function of time until merger for binaries at $z = 3$. The total mass M has been divided into three categories just as in Table 3.4. For rows marked "a," some very massive systems are excluded from the data. In those cases, the position and distance are very poorly constrained that far in advance of merger. In some cases, the binary is even out of band. For the row marked "b," all of the binaries of that mass range are either very poorly measured or completely out of band so far in advance of merger. 140

3.6 Typical ranges of sky position and distance measurement accuracy as a function of time until merger for binaries at $z = 5$. The total mass M has been divided into three categories just as in Table 3.4. For rows marked "a," some very massive systems are excluded from the data. In those cases, the position and distance are very poorly constrained that far in advance of merger. In some cases, the binary is even out of band. For the row marked "b," all of the binaries of that mass range are either very poorly measured or completely out of band so far in advance of merger. 141

4.1 General form of nonradial perturbations, split into odd and even parity and expanded as scalar, vector, and tensor spherical harmonics. . . . 163

4.2 Neutron star tidal disruption radius $r_{\rm tide}$, in units of black hole mass $m_{\rm BH}$, for various black hole-neutron star binaries and calculation methods. The binaries are parameterized by the neutron star's compactness $C = m_{\rm NS}/R$ and by the mass ratio $q = m_{\rm BH}/m_{\rm NS}$. "Full GR" results are calculated using numerical relativity [232]. "Quasi GR" results are calculated using a potential derived from the TOV equations [88]. "Newtonian" results use Newtonian self-gravity for the neutron star [254]. Finally, "perturbative GR" results are calculated using the methods of this chapter. 182

4.3 Additional number of GW cycles $\Delta N_{\rm cyc}$ caused by tidal effects for NS-NS binaries with $m_1 = m_2 = 1.4 M_\odot$. (The exception is GS1, which has a maximum mass less than $1.4 M_\odot$. In that case, each neutron star has the maximum mass.) 211

Chapter 1

Introduction

At the most basic level, this is a thesis about gravity. This begs the question: What is so interesting about gravity? Of the fundamental interactions of physics, gravity seems at first glance to be the simplest. Certainly, it is the one with which the nonscientist is most familiar; children learn that gravity "makes things fall" at an early age. It is also the interaction which physicists first sought to explain. Centuries before Maxwell completed his theory of electromagnetism, Newton proposed the inverse square law for gravity and used it to explain the observed orbits of the planets. It seemed that gravity was solved.

Indeed, Newton's theory stood unchallenged for hundreds of years. However, in 1859 it was discovered that the orbit of Mercury precessed more than could be accounted for by perturbations from the rest of the planets [164]. Even though the anomalous precession was only ~ 43 out of a total 5600 arcseconds per century[1], it called Newtonian gravity into question. In addition, Newton's theory was inconsistent with Einstein's 1905 theory of special relativity [75]. By 1916, Einstein had introduced the theory of general relativity (GR) to solve these problems [76]. GR proceeded to pass every observational test; it remains even today the accepted *classical* theory of gravity. This time, however, success lasted only until the 1920s, when the full-scale development of quantum mechanics introduced a new regime in which classical physics failed to describe reality. By the end of the 20th century, quantum field

[1]This is the modern value; in 1859, the discrepancy was measured to be ~ 38 arcseconds.

theories had been developed for electromagnetism, the weak force, and the strong force (the so-called "Standard Model") [190]. However, a quantum theory of gravity remains elusive. These experiences have taught that gravity is more complicated than one might first expect.

Even when the basic theory of gravity is well understood (in the relevant limit), it can still lead to complex behavior in practical situations. For example, while the Newtonian interaction between two bodies is relatively simple, the interaction between many bodies is extremely complicated. Much work has been done on the gravitational dynamics of astrophysical systems. These studies range from the relatively small—molecular clouds, globular clusters, and galaxies—to the relatively large—galaxy clusters and the large-scale structure of the universe (see the Millennium Simulation [230] for an example of the latter). In addition, Newtonian gravity can become more complicated in nonideal situations. For example, finite-sized objects feel the effects of tides. Even when described by a simple inverse square law, the observable effects of gravity can be quite complex.

The situation is even more interesting for general relativity. In GR, there are very few simple solutions to the field equations. Even the general two body problem is impossible to solve analytically, and it was only very recently solved numerically [201, 45, 19]. The simple solutions themselves come with unusual interpretations: a region of space from which nothing can escape, a universe that expands in time. More complicated solutions feature a great deal of rich behavior. This thesis describes two different astrophysical scenarios in which the effects of general relativity—in some cases, *additional* relativistic effects that were previously ignored—may have an important impact on the system. Both of these scenarios involve binaries in which each member is a compact object. The first is the coalescence of massive black holes in high-redshift galaxies. In these binaries, general relativity modifies the orbits from the standard Newtonian results. Especially interesting is the effect of spin on the orbital dynamics. The second scenario is the tidal interaction of a neutron star with a companion black hole or neutron star. Because neutron stars are such dense objects, their self-gravity, and its response to the tidal field, must be calculated using GR.

Interestingly, we even plan to observe these systems using an effect of general relativity: gravitational waves. These waves, generated by the motions of masses, can give additional information on systems which we observe through traditional astronomical methods and also discover systems which have not or cannot be observed with telescopes. Ultimately, we are interested in the imprint of the general relativistic effects on the gravitational wave signal and how well that imprint can be measured by a gravitational wave detector.

This introductory chapter serves to give a broad overview of the common background of these projects before delving into the specifics of each. It begins with a discussion of the theory of general relativity (Sec. 1.1), including a slew of definitions which will be useful throughout the thesis. Much of this material is based on [109, 48, 130]. Section 1.2 turns to astrophysics and briefly describes the nature of compact objects, focusing on neutron stars and black holes and the binaries they form. Next, Section 1.3 focuses on the primary observable, gravitational waves, including their basic nature and the status of current efforts to detect them. Finally, Section 1.4 gives a detailed outline of the rest of the thesis.

1.1 Einstein's theory of general relativity

1.1.1 The Equivalence Principle and its consequences

In 1905, Einstein proposed his theory of special relativity [75], which describes physics in inertial (and therefore "special") reference frames. The key feature introduced by special relativity is the combination of space and time, previously thought to be separate and unchanging concepts, into a four-dimensional structure known as spacetime. Theories such as electromagnetism can be described by fields defined at points of spacetime.

General relativity incorporates gravity into special relativity. The main concept of GR is that gravity is *not* treated as a field defined on top of spacetime. Instead, gravity is due to the geometry of the spacetime itself. In pure special relativity, the

geometry of the spacetime is flat, in a way which will be defined more precisely below. But in general relativity, spacetime is curved.

The description of gravity as geometry is a consequence of the universality of the gravitational interaction, codified in the so-called Equivalence Principle.[2] This principle, known since the experiments of Galileo, states that the inertial mass and gravitational mass of any object are the same. The inertial mass is the ratio between an applied force and the resulting acceleration in Newton's second law of motion. The gravitational mass is the constant of proportionality in Newton's universal law of gravitation. There is no reason why these two masses should be the same, and yet it seems as if they are. As a consequence, the mass (singular, since the two are identical) drops out of the calculation of the acceleration due to gravity. In simple terms, "all objects fall at the same rate." This is not true of other fundamental interactions; for example, two objects could have the same mass but different electromagnetic charges, causing them to accelerate at different rates (or even in different directions) when placed in an electric field.

A thought experiment shows the important consequences of this universality of free fall. An observer who is freely falling in a gravitational field (for example, an astronaut on the International Space Station) cannot tell the difference between his situation and being in an inertial frame far away from any source of gravitational field. All nearby objects (other astronauts, tools) are falling at the exact same rate. Einstein generalized this principle, stating that the observer can do *no experiment* that will detect the gravitational field. All the laws of physics will work the same as they did in the inertial frames of special relativity. It is logical, therefore, to *define* an observer in free fall to be an inertial observer and then define acceleration relative to that observer. In general relativity, particles are *not* considered to be accelerated by gravity. However, other forces can cause particles to accelerate off their inertial, freely falling paths.

There is a slight problem with this idea as we have currently summarized it. The

[2]Technically, this is the Weak Equivalence Principle; later, we generalize to the Einstein Equivalence Principle.

gravitational field in which the observer sits is not perfectly uniform. In space, the gravitational field points towards the center of the Earth. Therefore, an astronaut at one end of the ISS will be falling in a slightly different direction than one at the other end. In this way, the gravitational field *can* be detected. Nevertheless, in a small enough region of spacetime, no *local* experiment can determine the presence of the gravitational field. At each point, we can define a small "locally inertial" frame in which physics follows the laws of special relativity. However, we cannot define these frames on the global structure of spacetime. The deviation of particles due to the global structure of spacetime is manifestation of the *tidal force*, which is important for finite-sized bodies. In Chapter 4, we will investigate the effects of tidal forces on neutron stars. In general relativity, the tidal force is the true gravitational force.

1.1.2 Key features of general relativity

We now introduce the basic mathematical structure for general relativity. The most important concept is the metric. The metric, denoted $g_{\mu\nu}$, is a symmetric tensor defined at each point of spacetime.[3] The metric of special relativity's flat spacetime is given the special notation $\eta_{\mu\nu}$, where

$$\eta_{\mu\nu} = \text{diag}(-1, 1, 1, 1)\,. \tag{1.1}$$

The notation "diag" means that the object can be represented as a diagonal matrix with the following entries along the diagonal. Other authors choose a different "signature" so that (1.1) is multiplied by -1. The form of $\eta_{\mu\nu}$ is preserved by transformations between inertial frames (Lorentz transformations); in general, the metric $g_{\mu\nu}$ looks different when expressed using different coordinate systems. The inverse metric $g^{\mu\nu}$ is defined as

$$g^{\mu\lambda} g_{\lambda\nu} = \delta^{\mu}{}_{\nu}\,, \tag{1.2}$$

[3]In this thesis, Greek indices will always run from 0 to 3, with 0 representing the time coordinate and $1-3$ representing the three spatial directions. In most cases, Roman indices will run from $1-3$, representing just the three spatial directions. (The exception is in Chapter 4, when we briefly use Roman indices to refer only to the angles θ and ϕ.) We also equate the units of space and time by setting $c = 1$ throughout.

where $\delta^\mu_{\ \nu}$ is the Kronecker delta function: $\delta^\mu_{\ \nu} = 1$ if $\mu = \nu$ and 0 otherwise. Here, and throughout this thesis, we use the Einstein summation convention, in which repeated indices in superscript and subscript positions are summed over.

The metric determines the invariant distance between infinitesimally separated points. Consider two events in spacetime with coordinates x^μ and $x^\mu + dx^\mu$. The invariant distance between these events is given by

$$ds^2 = g_{\mu\nu} dx^\mu dx^\nu \,. \tag{1.3}$$

To find the distance along finite paths in spacetime, (1.3) must be integrated along the path. For two events which are *spacelike* separated, $ds^2 > 0$, and ds describes the "proper distance" between those events; that is, the distance measured in coordinates for which the events happen simultaneously. For two events which are *timelike* separated, $ds^2 < 0$, and $d\tau = \sqrt{-ds^2}$ is the "proper time" between those events, the time measured in coordinates at which both events happen at the same spatial location. Finally, *null* or *lightlike* events have $ds^2 = 0$.

Functionally, the metric is also useful for defining inner products of vectors and tensors. Consider two (four-)vectors $\vec{A}$ and $\vec{B}$, which have components A^μ and B^μ in some coordinate system.[4] The inner product between the two vectors is given by

$$\vec{A} \cdot \vec{B} \equiv g_{\mu\nu} A^\mu B^\nu \equiv A_\nu B^\nu = A^\mu B_\mu \,. \tag{1.4}$$

The last two equalities introduce the concept of *lowering* a vector with a superscript index (sometimes called a contravariant vector, or just a vector) to create the associated vector with a subscript index (sometimes called a covariant vector, a dual vector, or a one-form):

$$A_\nu = g_{\mu\nu} A^\mu \,. \tag{1.5}$$

Indices of tensors can be lowered using one copy of the metric for each index. Indices

[4]Throughout this thesis, we will use arrows (e.g. $\vec{A}$) to represent four-vectors and boldface (e.g. $\mathbf{S}$) to represent three-vectors. More commonly, though, we will just refer to a vector by its components (i.e., "the vector A^μ and the vector S^i").

can also be raised by the inverse metric $g^{\mu\nu}$.

Consider the series of events experienced by a massive (test) particle moving through spacetime. These events, which must be timelike separated, define the particle's trajectory or world line. (Massless particles like photons, on the other hand, follow a null trajectory.) The trajectory can be parameterized by the integrated proper time τ. The tangent to the world line is a vector known as the four-velocity u^μ:

$$u^\mu = \frac{dx^\mu}{d\tau}\,, \tag{1.6}$$

which is normalized as

$$u_\mu u^\mu = -1\,. \tag{1.7}$$

Particles in free fall (shielded from all nongravitational forces) move on special trajectories which maximize the proper time between two points. These special trajectories are known as "geodesics."[5] The geodesics can be described by the geodesic equation:

$$\frac{du^\mu}{d\tau} + \Gamma^\mu{}_{\rho\sigma} u^\rho u^\sigma = 0\,, \tag{1.8}$$

where

$$\Gamma^\sigma{}_{\mu\nu} = \frac{1}{2} g^{\sigma\rho}(\partial_\mu g_{\rho\nu} + \partial_\nu g_{\mu\rho} - \partial_\rho g_{\mu\nu}) \tag{1.9}$$

is known as the connection. Here $\partial_\mu \equiv \partial/\partial x^\mu$. In the flat spacetime of special relativity, $\Gamma^\sigma{}_{\mu\nu} = 0$ and geodesics are straight lines, as described by Newton's first law of motion. In curved spacetime, the geodesics are more complicated, but they can be viewed as generalized straight lines. This is the idea of the Equivalence Principle: Objects in free fall are not being accelerated away from straight lines. We have just altered the definition of "straight line" to reflect the geometry of spacetime.

One way to see this is to generalize the notion of partial derivative to curved spacetime. In curved spacetime, vectors defined at different points cannot be directly compared because each point has its own unique "tangent space" in which the vectors

[5] In more familiar curved *space*, such as the surface of the Earth, geodesics *minimize* the distance between two points. This is why airplanes fly on so-called great circles.

live. The *covariant derivative* corrects for this:

$$\nabla_\mu A^\nu = \partial_\mu A^\nu + \Gamma^\nu{}_{\mu\lambda} A^\lambda \,, \tag{1.10}$$

$$\nabla_\mu A_\nu = \partial_\mu A_\nu - \Gamma^\lambda{}_{\mu\nu} A_\lambda \,. \tag{1.11}$$

For tensors, the result is similar. The first term (the partial derivative) is the same, and then one term proportional to the connection is added (subtracted) for each superscript (subscript) index. Then (1.8) can be written as

$$u^\alpha \nabla_\alpha u^\mu \equiv \frac{D u^\mu}{d\tau} = 0 \,, \tag{1.12}$$

where we have defined $D/d\tau = u^\alpha \nabla_\alpha$ as the derivative along the geodesic. With an appropriate definition of the derivative in curved spacetime, it becomes much clearer that the particle is unaccelerated.

Consider a test gyroscope with spin four-vector $\vec{S}$. (In a frame where the gyroscope is at rest, $S^0 = 0$ and S^i is the standard three-dimensional spin $\mathbf{S}$.) In the flat spacetime of special relativity, the spin of the gyroscope is constant ($dS^\mu/d\tau = 0$) when no torques act upon it. Things become more complicated in curved spacetime. The result (1.12) inspires the notion that the spin is still constant, if we use the covariant derivative along the geodesic to define "constant":

$$\frac{D}{d\tau} S^\mu = 0 \,. \tag{1.13}$$

In flat spacetime, the connection is zero, and this statement says, as expected, that the spin does not precess. In curved spacetime, (1.13) says that while the spin is "parallel transported" along the geodesic (changed so that it always remains parallel to itself in the geometry of spacetime), it will precess with respect to the coordinates, which are fixed to the distant stars. This precession is known as geodetic precession.

The separation ξ^μ between two neighboring geodesics evolves to lowest order as

$$\frac{D^2 \xi^\mu}{d\tau^2} = R^\mu{}_{\nu\rho\sigma} u^\nu u^\rho \xi^\sigma \,, \tag{1.14}$$

where we have defined the *Riemann curvature tensor*

$$R^{\rho}{}_{\sigma\mu\nu} = \partial_\mu \Gamma^{\rho}{}_{\nu\sigma} - \partial_\nu \Gamma^{\rho}{}_{\mu\sigma} + \Gamma^{\rho}{}_{\mu\lambda}\Gamma^{\lambda}{}_{\nu\sigma} - \Gamma^{\rho}{}_{\nu\lambda}\Gamma^{\lambda}{}_{\mu\sigma} \,. \tag{1.15}$$

A spacetime which has zero curvature, like that of special relativity, is *flat*. In a flat spacetime, the geodesics will not deviate. This is similar to flat, Euclidean space, in which parallel lines remain parallel forever. In curved spacetime, however, the geodesics do deviate. This is the *tidal force*, the true force of gravity in general relativity. Note that, at lowest order, the tidal acceleration is proportional to the separation, just as it is in Newtonian gravity. In Chapter 4, we will use the Riemann tensor to calculate the tidal force on a neutron star.

We have now seen how the geometry of spacetime affects the motion of test particles in that spacetime. We must now briefly describe how the curvature of spacetime is determined in the first place. This is understood by analogy to Newtonian gravity, in which the mass of an object is the source of its gravity. Special relativity teaches us that mass and energy are equivalent, so energy should gravitate also. In general relativity, by analogy to electromagnetism, the *motion* of energy also creates gravity. The source of spacetime curvature in GR is known as the stress-energy tensor, $T^{\mu\nu}$. The elements of this tensor can be described as the flux of four-momentum p^μ in the x^ν direction. $T^{00} = T^{tt}$ is the local energy density. T^{ti} is the local energy flux, which is equal to T^{it}, the local momentum density. T^{ij} is the local momentum flux, including pressure terms on the diagonal and shear terms on the off-diagonal. We expect some analogy to energy conservation. In fact, the stress-energy is *locally* conserved:

$$\nabla_\mu T^{\mu\nu} = 0 \,. \tag{1.16}$$

However, in curved spacetime, there is no concept of global conservation of energy. The quantity $\nabla_\mu T^{\mu\nu}$ is a vector, and we cannot add up vectors defined at different points because they do not share tangent spaces in curved spacetime.

The equation relating the curvature of spacetime to the stress-energy tensor is

known as the Einstein equation. We first define the *Ricci tensor*:

$$R_{\mu\nu} = R^{\lambda}{}_{\mu\lambda\nu}\,. \tag{1.17}$$

The trace of this tensor is the *Ricci scalar*:

$$R = g^{\mu\nu}R_{\mu\nu} = R_{\mu}{}^{\mu}\,. \tag{1.18}$$

Putting these together, we form the *Einstein tensor*:

$$G_{\mu\nu} = R_{\mu\nu} - \frac{1}{2}Rg_{\mu\nu}\,. \tag{1.19}$$

This is the tensor which can be equated to the stress-energy tensor:

$$G_{\mu\nu} = 8\pi T_{\mu\nu}\,. \tag{1.20}$$

In the Einstein equation (1.20), and throughout this thesis, the gravitational constant $G = 1$. The system of units $G = c = 1$ is known as geometric units; in it, masses, lengths, and times all have the same dimension. A useful conversion factor is that $1\ M_{\odot} = 1.47$ km $= 4.92\ \mu$s.

The Einstein equation (or *equations*; there are six independent ones once the conservation principle is taken into account) is complicated and nonlinear. Therefore, it is extremely difficult to solve in general. Exact solutions only exist when the problem has a great deal of symmetry. For spherical symmetry and a vacuum ($T_{\mu\nu} = 0$), the solution is the Schwarzschild metric, which, as we shall see below, is the metric of a nonrotating black hole. The Schwarzschild metric is also the solution outside any spherical star. Inside, the metric is still spherically symmetric but more complicated; given a simple enough $T_{\mu\nu}$, modified stellar structure equations can be derived. Another exact solution is the Kerr metric, which describes an axisymmetric, rotating black hole (or the spacetime outside an axisymmetric, rotating body). Finally, by making assumptions of homogeneity and isotropy, a simple metric can be written

down for the entire expanding universe. Einstein's equations reduce to the Friedmann equations for the scale factor a.

For less symmetric situations, Einstein's equations cannot be solved analytically.[6] The obvious solution is to solve them numerically; however, this is easier said than done. GR naturally treats space and time as the same, while numerical simulations work by evolving functions through time. The best slicing of spacetime into space and time is not easy to determine. Other problems also exist, including boundary conditions for black holes, and how to start the abstracted simulation with realistic, astrophysical initial data. For years, problems like these stymied the field of numerical relativity, although great breakthroughs have been made in the past few years [201, 45, 19].

Even though numerical relativity has recently shown great promise, it is still a developing field. Furthermore, full numerical simulations can be computationally expensive. For these reasons, approximation is still a critical tool in general relativity research. To start, many interesting results can be derived in the weak field, or linearized gravity, limit, in which we take the metric to be that of flat space plus a small perturbation:

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu} \,. \tag{1.21}$$

When working in this limit, we can raise and lower indices of first order quantities using only the background metric $\eta_{\mu\nu}$. In those cases (mainly in the following section on gravitational waves), we are less careful about index placement. Weak field theory is appropriate to describe Newtonian gravity, as well as various modifications from it, like the correct expression for the bending of light by the Sun. When *dynamical* effects are taken into account, linearized gravity produces gravitational waves, which are introduced below.

Post-Newtonian theory [34] is an extension of weak field theory; it can be considered an expansion in the strength of the gravitational field and the internal velocities of the source. Numerical relativity simulations have shown that post-Newtonian re-

[6] Here, "analytically" means that they can, at least, be reduced to a set of easily integrable ordinary differential equations.

sults are remarkably good at describing relativistic compact binaries over the duration of the inspiral [186]. We use post-Newtonian results extensively in Chapters 2 and 3.

Perturbation theory is also an extension of weak field theory. Instead of adding a small perturbation to flat space, a small perturbation is added to a more general background metric. For example, black hole perturbation theory is useful in describing the final "ringdown" of a hole created in a binary merger [146]. In Chapter 4, we use perturbation theory to study the tidal distortion of neutron stars.

1.2 Compact objects and binaries

A compact object is one which has an unusually small radius R for its mass m. They come in three varieties; in order of increasing m/R, these are: white dwarfs, neutron stars, and black holes. In this section, we discuss the properties of these objects and the binaries they form.

1.2.1 White dwarfs

White dwarfs (WDs) are the remnants of low- to medium-mass stars [47]. After these stars finish nuclear burning, they expel their outer layers as a planetary nebula. The core—usually carbon-oxygen, though helium and oxygen-neon-magnesium cases exist—is left behind as a white dwarf. These stars typically have masses of $\sim 0.6\ M_\odot$ but small radii of only $\sim 0.01\ R_\odot$ (approximately the radius of the Earth). Gas and radiation pressure are inadequate to support such stars. Instead, they are supported by electron degeneracy pressure, implying a maximum mass of $\sim 1.44\ M_\odot$, the Chandrasekhar limit [49].

When a white dwarf approaches the Chandrasekhar limit, due to accretion [253] or (less likely) merger with another dwarf [134, 252], it will begin a runaway thermonuclear reaction that destroys the star. This process is the favored mechanism for so-called "Type Ia supernovae" [126]. Because all Type Ia supernovae occur under similar conditions, they are very useful as "standard candles" [192, 207, 251]. The distance to a distant supernova can be obtained by comparing its observed magni-

tude to that of a closer supernova whose distance is already known. By combining distance measurements with redshift measurements, the cosmography of the universe can be established. In 1998, Type Ia supernovae measurements were used to show that the universe is accelerating [206, 189]. This unexpected result is attributed to "dark energy," an unexplained phenomenon which makes up $\sim 70 - 75\%$ of the universe by energy content [118]. However, there are many systematic errors in the conversion of luminosity to distance which make precise measurements of dark energy difficult. While researchers are attempting to better understand the white dwarf explosion mechanism, it would be nice to have another, independent method to measure cosmological parameters. One possibility is discussed in Chapter 3.

The role of white dwarfs in this thesis is small. In Chapter 2, we shall see that white dwarf binaries are an interesting source of low-frequency gravitational waves—one that also becomes a liability if we are trying to look at other sources. In Chapter 4, we will use a white dwarf to test tidal perturbation calculations. Since white dwarfs are much less compact than neutron stars, their self-gravity is well approximated by Newtonian methods.

1.2.2 Neutron stars

Neutron stars (NSs) are compact objects which are composed primarily of tightly packed neutrons. Just two years after the discovery of the neutron, Walter Baade and Fritz Zwicky predicted [17] that neutron stars would be the the end stage of medium- to high-mass stars which end their lives with core-collapse (Types Ib, Ic, and II) supernovae.

Neutron stars are usually observed as pulsars, sources of radiation that repeat, or pulse, extremely regularly. The first pulsar was observed in 1967 by Jocelyn Bell and Antony Hewish [115] and was later determined to be a rotating neutron star [103]. In 1974 a so-called "binary pulsar," consisting of a pulsar in orbit around another neutron star, was found by Hulse and Taylor [133]. By measuring the arrival time of the pulses, information about the binary orbit can be obtained. Since the system is strongly relativistic, it serves as a "laboratory" for testing many of the predictions of

general relativity [234, 235]. As we shall see in the next section, the pulsar provided the first indirect evidence for gravitational waves.

Neutron stars have masses $\sim 1-2\ M_\odot$ and radii ~ 10 km, with an average density greater than nuclear densities. The nature of matter at such high densities is poorly understood [161, 162, 163]. Many possible equations of state, or relations between pressure and density, have been calculated by various authors. (See Chapter 4 for a large list.) In some of these possibilities, the core of the star contains more exotic matter, such as pions, hyperons, kaons, or free quarks, in addition to nucleons and leptons. In the extreme case, the entire star may spontaneously convert to strange quark matter, making the term "neutron star" a bit of a misnomer. Astronomical measurements of neutron stars can help distinguish between these different possibilities. Traditionally, these measurements consist of the mass and radius of the neutron star. In Chapter 4, we review these ideas and then investigate what additional information we can learn from the neutron star's response to tidal forces.

Because neutron stars are so dense ($m/R \sim 0.1 - 0.2$), any calculation of the structure (e.g., for a given equation of state) must include general relativistic effects. For the approximations of a spherical star made of a perfect fluid (one without shear stresses, viscosity, or heat conduction), Einstein's equations can be reduced to a set of stellar structure equations, known as the Tolman-Oppenheimer-Volkoff equations [244, 184]. We briefly derive these (well-known) equations in Chapter 4. We also add tidal distortions to the stars using perturbation theory.

1.2.3 Black holes

A black hole (BH) is a region of spacetime in which gravity is so strong that even light cannot escape [109, 48]. The boundary of a black hole which marks the point of no escape is known as the event horizon, or simply the horizon. All timelike and null trajectories which begin inside the horizon remain inside.

Stationary black holes have the interesting property that they are completely described by their mass, spin, and charge. The black hole is said to have "no hair," or other distinguishing features, including any information about the mass that may have

formed it in the first place. Because of this property, the spacetime of a stationary black hole is easy to describe. The simplest holes, those without spin or charge, are described by the Schwarzschild solution to the Einstein equations:

$$ds^2 = -\left(1 - \frac{2m}{r}\right) dt^2 + \frac{dr^2}{1 - 2m/r} + r^2(d\theta^2 + \sin^2\theta d\phi^2)\,, \tag{1.22}$$

Clearly, this metric is poorly behaved at $r = 2m$. More careful investigation shows that this is the location of the horizon. (The black hole can then be said to have a size $R = 2m$ or $R \sim m$.) The spacetime itself is well behaved at the horizon (i.e., the curvature is finite); the singularity is due purely to the choice of so-called "Schwarzschild coordinates."[7] At $r = 0$, on the other hand, the spacetime is badly behaved in *all* coordinate systems. This is a true curvature singularity, a point where the classical GR description is not adequate. Luckily, this singularity is hidden from the outside universe by the horizon, so the details at that point do not matter for calculations involving black holes.[8]

More general is the Kerr metric, which describes rotating black holes:

$$\begin{aligned} ds^2 = &-\left(1 - \frac{2mr}{\Sigma}\right) dt^2 - \frac{4amr\sin^2\theta}{\Sigma} dt d\phi + \frac{\Sigma}{\Delta} dr^2 + \Sigma d\theta^2 \\ &+ \left(r^2 + a^2 + \frac{2mra^2\sin^2\theta}{\Sigma}\right) \sin^2\theta d\phi^2\,, \end{aligned} \tag{1.23}$$

where

$$\Delta = r^2 - 2mr + a^2\,, \tag{1.24}$$

$$\Sigma = r^2 + a^2\cos^2\theta\,, \tag{1.25}$$

and

$$a = \frac{|\mathbf{S}|}{m} \tag{1.26}$$

is a measurement of the spin of the black hole, which is restricted to $a \leq m$. We shall

[7]Note that in these coordinates, r labels surfaces of area $4\pi r^2$ but does not describe the proper distance to that surface.

[8]The singularity *can* cause trouble in numerical relativity if one is not careful.

also have cause to use a dimensionless spin parameter $\chi = a/m = |\mathbf{S}|/m^2$. The metric (1.23) is written is so-called "Boyer-Lindquist" coordinates, in which the horizon is located at $r = m + \sqrt{m^2 - a^2}$. For $a = 0$, this reduces to the Schwarzschild result, but for larger a, the horizon moves inward as far as $r = m$.

The $g_{t\phi}$ term in the Kerr metric is a *gravitomagnetic* effect. The rotation of the star "drags" inertial frames with it. A test gyroscope sitting in the Kerr spacetime will therefore precess relative to the distant stars, even if it is held at fixed coordinates. This precession, called Lense-Thirring precession, acts in concert with the geodetic precession mentioned earlier. These precession effects, generalized to extended, non-test bodies, are the main focus of Chapter 2.

Additional black hole solutions exist which add charge to the Schwarzschild and Kerr solutions, but since we expect the charge of any real black holes to be quickly neutralized by nearby matter, we can take the Kerr solution as the most general black hole solution.

As presented so far, black holes represent particular solutions to Einstein's equations and not actual astrophysical objects. However, there is convincing evidence that black holes do actually exist in nature. Stellar-mass black holes ($\sim 3 - 100\ M_\odot$) are believed to be the final state of the most massive stars [112]. In some scenarios, a neutron star formed during a supernova accretes enough infalling matter to surpass its (equation-of-state-dependent) maximum mass and then collapses into a black hole. In other scenarios, the mass of the progenitor star is large enough that it collapses directly to a black hole. Since black holes cannot be seen directly, they can only be observed indirectly in X-ray binaries. An X-ray binary consists of a main sequence star and a compact object. Matter from the normal star accretes onto the compact object, releasing radiation in the form of X-rays. From the set of all X-ray binaries, black hole candidates can be identified as those compact objects which have a mass greater than the greatest possible maximum neutron star mass ($\sim 3\ M_\odot$). In addition, neutron stars have surfaces on which the buildup of material can lead to thermonuclear bursts. If these bursts are not seen, especially in stars much more massive than $\sim 3\ M_\odot$, it is logical to conclude that the compact object in question

has no surface and is a black hole.

At the other end of the mass scale are massive ($\sim 10^4 - 10^7\ M_\odot$) and supermassive ($\sim 10^7 - 10^9\ M_\odot$) black holes.[9] The radio source Sagitarrius A* at the center of the Galaxy is believed to be a massive black hole. Measurements of nearby stellar orbits imply the existence of a mass $\sim 4 \times 10^6 M_\odot$, localized in an extremely small area [74, 100]. In fact, massive black holes seem to be ubiquitous in the local universe: It appears that all galaxies with central bulges contain black holes [148, 168] whose masses are strongly correlated with the properties of the bulge [87, 99]. Meanwhile, quasars, which are powered by accretion onto a supermassive black hole, have been observed at high redshift ($z \sim 6$) [83]. The formation of these massive and supermassive black holes is still poorly understood. A popular model is the collapse of massive Population III (i.e., first generation) stars, followed by accretion. Another model is the direct collapse of a large gas cloud, giving the black hole a large mass at formation [220].

In between these limits are intermediate-mass black holes ($\sim 10^2 - 10^4 M_\odot$). Evidence for these holes is the weakest of all three cases. It is possible that the brightest ultra-luminous X-ray sources (ULXs) in nearby galaxies are actually accreting intermediate-mass black holes [53]. Such holes may be formed by the collisions of massive stars in globular clusters [195]. While intermediate-mass black holes are a subject of much study, we shall not discuss them further in this thesis.

1.2.4 Compact binaries

The subject of this thesis is compact binaries; unlike the X-ray binaries described earlier, these are binaries in which both members are compact objects. We focus on *relativistic* binaries in which each member is either a neutron star or black hole, although as we have mentioned, white dwarf binaries are a significant source of noise in our analysis of Chapters 2 and 3.

Binaries involving neutron stars and/or *stellar-mass* black holes may begin as

[9] This distinction is not well defined and often not defined at all. We separate the two ranges based on the sensitivity of LISA to massive black hole binaries.

normal binaries of main-sequence stars, where each has enough mass to eventually go supernova. The more massive star will evolve first into a giant, transfer mass to its companion (possibly putting it over the supernova limit if it was not before), and then finally explode, leaving behind a compact remnant. Depending on how much mass is lost in the supernova, the binary may remain bound or become unbound. If it stays bound, the second star will eventually follow the same path as the first. When it transfers mass to the compact star, it can be seen as a high-mass X-ray binary. (The fact that these binaries are observed means that the scenario to this point must be plausible.) Eventually, the compact object will be encased by the envelope of the giant. This is known as the "common envelope" phase of evolution. Drag forces take energy from the binary orbit and eventually expel the envelope. The giant then can go supernova as well; if the binary survives this explosion, a NS-NS, NS-BH, or BH-BH binary is formed [233].

Stellar-mass binaries can also be formed by another mechanism. In globular clusters, massive objects (such as compact objects or their progenitors) tend to sink to the center, where they can interact and form binaries. This process is especially effective at producing BH-BH binaries [211]. The total compact binary formation rate includes contributions from both of these processes.

As we have seen, NS-NS binaries have been definitively observed to exist. The most famous example is the Hulse-Taylor binary pulsar, but several other binary pulsars have also been discovered, including the *double* binary pulsar PSR J0737-3039A & B [166]. NS-BH binaries have not been observed. However, it is believed that compact binary mergers are the source of many short, hard gamma-ray bursts [181]. Since matter is required to produce the gamma rays, these binaries must be NS-NS or NS-BH. Finally, there is no evidence, direct or indirect, for the existence of stellar-mass BH-BH binaries. These objects are completely dark and can only be detected by the gravitational radiation they emit. In this thesis, we will focus only on the binaries containing neutron stars, NS-NS and NS-BH.

Binaries involving *massive* (and supermassive) black holes come in two flavors. The first is the capture of a "small" compact object (WD, NS, stellar-mass BH) by

a large black hole [228]. We shall see in the next section that these binaries make excellent sources of gravitational waves. However, they shall not interest us in the rest of this thesis. The other possibility is a massive black hole binary (MBHB). As we have mentioned, large black holes are observed in the cores of nearly all nearby galaxies. In addition, hierarchical structure formation teaches us that these galaxies assembled over cosmic history through the repeated coalescence of the dark matter halos in which they reside [69]. Taken together, these facts suggest that MBHBs should be relatively common, especially at high redshift.

The formation scenario for MBHBs was first described in [24]. When two galaxies merge, their massive black holes sink to the center of the new galaxy due to dynamical friction. At a separation of $\sim$ 1 parsec, they form a binary. The binary hardens (shrinks its radius) by interacting with nearby stars. In these three-body interactions, the stars take energy from the binary, causing them to be ejected from the galactic center. When the holes get close enough ($\sim 10^{-3}$ pc), gravitational radiation becomes effective at shrinking the orbit, and the holes will eventually coalesce. Unfortunately, many N-body simulations have had difficulty getting the holes down to this small separation. The problem is that only a limited number of stars are on so-called "loss cone" orbits which can interact with the binary to remove energy. When these orbits are depleted, the binary stalls at $\sim$ 1 pc [176]. Recently, there have been solutions to this "last parsec" problem, including the effects of gas on the orbit [11, 78, 79] or the use of a triaxial potential [27, 28].

Observations have provided some evidence for this picture. First, galaxies have been seen to have two MBHs, separated by a few kiloparsecs [147, 32]. While these are too widely separated to be considered a binary, they may be the precursor to one. More convincing is the observation of two MBHs separated by $\sim$ 7 pc [209]. In addition, the quasar OJ 287 produces quasi-periodic bursts which may be explained by a binary [249]; if this interpretation is correct, the holes are only separated by $\sim$ 0.05 pc, providing the best evidence yet for the existence of massive black hole binaries.

1.3 Gravitational waves

In the previous section, we discussed several different varieties of relativistic compact binaries. Only one of these varieties has actually been observed (NS-NS), and then only in a widely separated regime. We may be observing merging NS-NS and NS-BH systems in gamma-ray bursts, but many details about this picture remain uncertain. Stellar-mass BH-BH systems have not been observed, and only tantalizing hints exist for their massive cousins. However, the situation may change in the next decade, as all of these systems are strong emitters of gravitational waves.

In electromagnetism, waves are generated by accelerating charges. We expect that in general relativity, gravitational waves should be produced by accelerations of mass-energy.[10] However, there are fundamental differences between electromagnetic (EM) waves and gravitational waves (GWs) [128, 97]:

- EM waves are the oscillations of electric and magnetic fields defined on spacetime, while GWs are the oscillations of spacetime itself. It is not immediately clear how to separate these oscillations from the rest of spacetime and define a GW. To do so requires a separation of length scales; for a review, see [97].

- EM waves are created, at lowest order, by the time-changing charge *dipole* moment of the source. The monopole moment cannot change due to charge conservation. For GWs, we also expect the monopole moment to be unchanging due to mass-energy conservation. However, the mass-energy and mass-energy-current dipole moments also cannot change due to momentum and angular momentum conservation. Therefore, the leading order gravitational radiation comes from time-changing *quadrupole* moments.

- A linearly polarized EM wave accelerates test charges back and forth in the direction of the electric field. GWs act in a *tidal*, quadrupolar fashion: A linearly polarized GW stretches spacetime in one direction and squeezes it in the orthogonal direction (and then vice-versa as the wave passes). This means

[10] Here we are using the term "acceleration" loosely, since GWs can be generated by purely gravitational, and thus "unaccelerated" motion.

that the two polarizations of GWs are separated by 45 degrees, not 90 as for EM waves. This is consistent with a quantum-mechanical spin-2 particle (the "graviton"), compared to the spin-1 photon.

- EM waves are usually observed as an energy flux, which drops off as $1/r^2$ with distance. The GW strain, introduced below, falls off as $1/r$, meaning that improvements in sensitivity translate linearly to improvements in observable distances.

- EM waves interact strongly with matter, while GWs do not. This is both a blessing and curse in each case. Strong interactions make EM waves easier to detect, but it also means that they are often scattered or absorbed between the source and the observer. (Of course, these processes often tell us a great deal about the space in between!) Weak interactions mean that GWs can be detected from great distances and behind intervening matter. But it also means that once the GWs arrive on Earth, they are extremely difficult to detect.

- EM waves are generated by various small particles, such as accelerating electrons or atoms emitting line radiation. We see the incoherent superposition of all these radiative processes. GWs are coherently produced by the bulk motions of the system, such as the motion of two stars in orbit around each other. This phase coherence helps us detect weak GW signals by matching them to various templates over a long integration time ("matched filtering").

- The wavelengths of EM waves are usually small enough to allow imaging of a system, while GWs have wavelengths which are much too large for imaging. In this way, GWs are more analogous to sound waves than to EM waves. Later in this paper, we will refer to certain GW sources as "standard sirens,"[11] the GW analogy to EM "standard candles" such as Type Ia supernovae.

- Similarly, while EM waves are detected by focusing on a specific point on the sky, GW detectors "hear" sources from the entire sky. Unlike an EM source,

[11]This term was coined by Sterl Phinney and Sean Carroll.

which is localized by default when it is detected, GW sources have relatively poorly determined sky positions. Chapters 2 and 3 of this thesis discuss the localization potential for one particular GW source, massive black hole binaries.

In the rest of this section (based on elements of [128, 48, 97, 130]), we discuss the basic mathematics of gravitational waves, potential astrophysical sources for the waves, and the ongoing efforts to detect them.

1.3.1 Basic theory of gravitational waves

We begin by considering weak-field gravity, with metric perturbation $h_{\mu\nu}$. In Chapter 4, we present the perturbations to the connection and the various curvature tensors for perturbations to a general background, so we will not bother to present the (simpler) flat background case here. Instead, we just give the linearized Einstein tensor:

$$G_{\mu\nu} = \frac{1}{2}\partial_\mu\partial^\lambda h_{\lambda\nu} + \partial_\nu\partial^\lambda h_{\mu\lambda} - \partial_\mu\partial_\nu h - \Box h_{\mu\nu} + \eta_{\mu\nu}\Box h - \eta_{\mu\nu}\partial^\lambda\partial^\rho h_{\lambda\rho}\,, \tag{1.27}$$

where $h \equiv \eta^{\mu\nu}h_{\mu\nu}$ is the trace of $h_{\mu\nu}$ and $\Box \equiv \eta^{\mu\nu}\partial_\mu\partial_\nu$. If we define the "trace-reversed" metric perturbation,

$$\bar{h}_{\mu\nu} \equiv h_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu}h\,, \tag{1.28}$$

this equation can be simplified. More simplifications can be made by making an infinitesimal change to the coordinate system, thus changing the form of $h_{\mu\nu}$. This is the expression of *gauge freedom* in general relativity. If we change coordinates so that $x^\mu \to x^\mu + \xi^\mu$, then the perturbation changes as

$$h_{\mu\nu} \to h_{\mu\nu} - \partial_\mu\xi_\nu - \partial_\nu\xi_\mu\,. \tag{1.29}$$

(Note that in Chapter 4, for a general background, partial derivatives in this expression are replaced by covariant derivatives.) Changing gauge leaves the Riemann tensor unchanged, leading to the interpretation of the metric perturbation as a "grav-

itational potential" and the curvature as the true "gravitational field." If we choose the so-called Lorenz gauge, $\partial^\mu \bar{h}_{\mu\nu} = 0$, we finally obtain the linearized Einstein equation

$$\Box \bar{h}_{\mu\nu} = -16\pi T_{\mu\nu} \,. \tag{1.30}$$

When solving for gravitational wave solutions, it turns out that we only need to worry about the spatial, transverse, and traceless components of the metric, denoted h_{ij}^{TT}. All of the other components of the metric represent gauge degrees of freedom, or physical, but nonradiative, degrees of freedom related to the presence of nearby matter [48, 97]. Transverse means that h_{ij}^{TT} is orthogonal to the direction of the wave's propagation:

$$\partial_i h_{ij}^{TT} = 0 \,, \tag{1.31}$$

just like the EM fields of an EM wave. Traceless means that

$$\delta_{ij} h_{ij}^{TT} = 0 \,. \tag{1.32}$$

To convert h_{ij} to h_{ij}^{TT}, use the formula:

$$h_{ij}^{TT} = h_{k\ell} \left(P_{ki} P_{\ell j} - \frac{1}{2} P_{k\ell} P_{ij} \right) \,, \tag{1.33}$$

where

$$P_{ij} = \delta_{ij} - n_i n_j \tag{1.34}$$

is a projection tensor for waves propagating along the unit vector $\hat{\mathbf{n}}$. Since the final result is traceless, $\bar{h}_{ij}^{TT} = h_{ij}^{TT}$.

In vacuum, the solutions to (1.30) are just plane waves, $h_{ij}^{TT} = C_{ij} \exp(i k_\mu x^\mu)$. The vector k_μ is null ($k_\mu k^\mu = 0$, so GWs travel at the speed of light) and orthogonal to C_{ij} on all spatial indices. By virtue of symmetry and tracelessness, C_{ij} can only have two unique components; these are the two polarizations ("plus" and "cross") of the wave. For a wave moving in the z (i.e., x^3) direction, we have

$$h_{ij}^{TT} = \begin{bmatrix} h_+ & h_\times & 0 \\ h_\times & -h_+ & 0 \\ 0 & 0 & 0 \end{bmatrix} . \tag{1.35}$$

If the wave is moving in an arbitrary direction, the result is more complicated; see Chapter 2 for details.

When a matter source is present, the trace-reversed metric perturbation is given by:

$$\bar{h}_{\mu\nu}(\mathbf{x}, t) = 4 \int \frac{T_{\mu\nu}(\mathbf{x}', t - |\mathbf{x} - \mathbf{x}'|)}{|\mathbf{x} - \mathbf{x}'|} d^3x' . \tag{1.36}$$

If we assume a distant, slow-moving source, then $|\mathbf{x} - \mathbf{x}'| \approx r$, and we can derive

$$h_{ij}^{TT} = \frac{2}{r} \frac{d^2 I_{k\ell}(t - r)}{dt^2} \left(P_{ik} P_{j\ell} - \frac{1}{2} P_{k\ell} P_{ij} \right) , \tag{1.37}$$

where

$$I_{ij}(t) = \int x^{i'} x^{j'} T_{tt}(\mathbf{x}', t) d^3x' \tag{1.38}$$

is the mass-energy quadrupole moment of the source. Better approximations show that this is only the first term in a multipolar expansion of the GWs, including octupole and higher mass-energy moments as well as quadrupole and higher mass-energy-current moments. As stated earlier, the fundamental radiation is quadrupolar; conservation laws prevent the appearance of dipole terms. For a binary source, this implies that the fundamental GW phase and frequency are equal to twice the orbital phase and frequency.

If we were to return to nongeometric units, we would find that the wave amplitude is scaled by G/c^4, a very small number. Since, as we shall see, the wave amplitude translates directly into an observable *strain*, this is indicative of how difficult GWs are to measure. We can only hope to detect sources with a large quadrupole moment which changes rapidly. Only the most extreme objects in the universe make detectable GW sources.

Even though the waves are difficult to detect, they still contain a great deal of

energy. Defining the energy in a GW requires going to another order in perturbation theory. The calculation can be found in [48, 97, 130]; the result is the Isaacson stress-energy tensor:

$$T^{\mathrm{GW}}_{\mu\nu} = \frac{1}{32\pi}\langle(\partial_\mu h^{TT}_{\rho\sigma})(\partial_\nu h^{\rho\sigma}_{TT})\rangle\,, \tag{1.39}$$

where the angle brackets mean to average over several wavelengths. (For a curved background, the partial derivatives are replaced with covariant derivatives on the background.) From the definition of the stress-energy tensor, the energy flux is given by $T^{\mathrm{GW}}_{0k}n^k$. Integrating over all possible directions n^k, the total luminosity L_{GW} emitted from a GW source is

$$L_{\mathrm{GW}} = \frac{1}{5}\left\langle \frac{d^3\mathcal{I}_{ij}}{dt^3}\frac{d^3\mathcal{I}_{ij}}{dt^3}\right\rangle\,, \tag{1.40}$$

where $\mathcal{I}_{ij} \equiv I_{ij} - \delta_{ij}I_{kk}/3$ is the reduced (traceless) quadrupole moment. Although this expression carries a factor of G/c^5 in nongeometric units, the squaring of the quadrupole moment helps to overcome this and produce incredibly large energies. While gravitational waves have not yet been detected directly, their existence has been confirmed *indirectly* by using this energy loss formula. Specifically, the orbit of the Hulse-Taylor binary pulsar was measured over the course of many years. The orbit was found to decay in a manner consistent with the energy loss predicted by (1.40) [234, 235].

1.3.2 Gravitational wave sources and detection

Gravitational wave sources can be divided into different frequency bands based on the detection method which is appropriate. As defined by Hughes [128], these are: high frequencies (1 Hz $\leq f \leq 10^4$ Hz), low frequencies (10^{-5} Hz $\leq f \leq 1$ Hz), very low frequencies (10^{-9} Hz $\leq f \leq 10^{-7}$ Hz), and ultra low frequencies (10^{-18} Hz $\leq f \leq 10^{-13}$ Hz). The high-frequency band contains a variety of potential GW sources, including the core collapse of massive stars in supernovae (but only if they have a nonspherical component), periodic sources like rotating neutron stars (if nonaxisym-

metric), and the focus of this thesis, compact binaries. In this frequency range, they are composed of two stellar-mass compact objects (NS-NS, NS-BH, and BH-BH). Unlike the Hulse-Taylor binary pulsar, which has a relatively long orbital period (7.75 h), these binaries are near the end of the orbital decay process. The low-frequency band contains their final orbits and merger.

These sources are currently being searched for by ground-based *interferometric detectors*. These detectors operate on the principle that a passing GW stretches and squeezes the space between two test masses, creating a time-dependent length change. For example,

$$\frac{\delta L_x(t)}{L} \approx \frac{1}{2} h_{xx}^{TT}(t, z=0) = \frac{1}{2} h_+(t, 0) \tag{1.41}$$

for two masses on the x-axis and a wave moving along the z direction. Similarly,

$$\frac{\delta L_y(t)}{L} \approx -\frac{1}{2} h_+(t, z=0) \tag{1.42}$$

for two masses on the y-axis. The x arm is stretched while the y arm is squeezed, and vice-versa; this is the quadrupolar response mentioned above. Putting these two length changes together, we find the *strain*

$$h(t) = \frac{\delta L_x(t) - \delta L_y(t)}{L} = h_+(t)\,, \tag{1.43}$$

which can be measured using laser interferometry. (A detector rotated by 45 degrees with respect to this one can measure the second polarization, $h_\times$. In reality, the detector will measure a weighted combination of the two polarizations depending on the direction of the source; see Chapter 2 for more details.) This technique is quite difficult to implement because the strain $h \lesssim 10^{-21}$. In addition to this small signal, ground-based interferometers must contend with various noise sources, including seismic noise, thermal excitations, and photon shot noise [1].

A network of these interferometers has been set up around the globe. For example, the Laser Interferometer Gravitational Wave Observatory (LIGO) consists of three Michelson interferometers, two ($L = 4$ km and $L = 2$ km) in Hanford, Washington,

and one ($L = 4$ km) in Livingston, Louisiana [1]. The Virgo detector ($L = 3$ km) is located in Cascina, Italy [39]. Smaller detectors also exist: GEO600 ($L = 600$ m) is located near Hannover, Germany [257], and TAMA300 ($L = 300$ m) is located in Mitaka, Japan [9]. LIGO has achieved its initial design sensitivity, which can detect a binary of two 1.4 $M_\odot$ neutron stars to ~ 15 Mpc [138]. Recent estimates predict detection rates of $\sim 0.002 - 0.005$ yr^{-1} for NS-NS binaries, $\sim 0.0002 - 0.0003$ yr^{-1} for NS-BH binaries, and $\sim 0.0003 - 0.0006$ yr^{-1} for BH-BH binaries [25]. Accounting for binaries formed in clusters can increase the BH-BH rate significantly [211]. LIGO's fifth science run (S5) ran from 2005-2007, obtaining a year of triple-coincident data. During this time, Virgo and GEO sometimes also ran coincidentally with the LIGO detectors [138]. No GWs have yet been reported, but upper limits have been obtained (including an intriguing limit on whether a short gamma-ray burst was indeed caused by a compact binary merger [54]).

A sixth science run will soon begin on so-called "Enhanced LIGO." Modifications have been made to the detectors which will increase the sensitivity by a factor of ~ 2 and thus the volume of observable space by a factor of ~ 8. (However, this improvement only holds above ~ 100 Hz [173].) When S6 is complete, construction will begin on "Advanced LIGO," which will increase the sensitivity (over initial LIGO) by a factor of ~ 10, for a factor of ~ 1000 improvement in volume [138]. Recent estimates predict that Advanced LIGO will detect $\sim 11 - 19$ NS-NS binaries, $\sim 0.68 - 1.3$ NS-BH binaries, and $\sim 1.1 - 2.5$ BH-BH binaries per year [25]; cluster formation of BH-BH binaries could bring the total merger rate up to $\sim 25 - 300$ yr^{-1} [211]. If these event rate estimates are correct, Advanced LIGO is virtually guaranteed to detect stellar-mass compact binaries. Chapter 4 discusses tidal effects which could have an impact on the GW signals from NS-NS and NS-BH systems.

The low-frequency band also contains a number of interesting compact binary sources. The first is the nearby population of white dwarf-white dwarf binaries [183, 84]. Many millions of these sources exist; while many of the signals can conceivably be individually resolved, the majority will remain unresolved. This makes them an important source of "confusion" noise for people looking for other sources. We include

an estimate of this confusion noise in our work of Chapters 2 and 3.

The other main low-frequency sources involve massive black holes. One interesting source is an "extreme mass ratio inspiral," created when a large black hole captures a small compact object. These events could provide a detailed probe of the spacetime of the large black hole, helping to verify the Kerr metric observationally [129]. In this thesis, however, we shall be concerned with a different source, binaries involving two massive black holes (with total mass $10^4 - 10^7 M_\odot$). Low-frequency GW detectors will be able to observe these binaries for days to years as they inspiral toward a final merger and "ringdown" into a single black hole.

The Laser Interferometer Space Antenna (LISA) is the proposed NASA-ESA mission to detect low-frequency gravitational waves [26]. Seismic noise makes it impossible to detect these waves on Earth, so LISA will be a space-based detector. It will consist of three spacecraft orbiting the Sun in a triangular configuration at 1 AU, trailing the Earth in its orbit. The test masses are shielded from all nongravitational forces so that they follow geodesic orbits. Unlike the ground-based detectors, LISA is fundamentally a "software interferometer"; signals are sent to neighboring spacecraft and compared to a local reference, but not reflected back. The phase differences at each spacecraft can later be combined to produce an interferometer signal. LISA's arm length is $\sim 5 \times 10^6$ km, so its length resolution does not need to be as precise as LIGO's. Noise sources include spurious accelerations of the test masses, shot noise, and optical-path measurement errors [159]. At high frequencies, the sensitivity is inhibited because the GW wavelength becomes shorter than the constellation arm length. This also makes the response function a bit more complicated than the simple "strain" picture described above [61, 210]. In addition, because the distance between spacecraft is unequal and time-varying, more complicated interferometry techniques must be used to eliminate laser phase noise [13].

LISA can detect massive black hole binaries out to very high redshift ($z > 10$) with reasonable signal-to-noise (SNR $\sim$ 100) and "nearby" binaries ($z \sim 1 - 5$) with outstanding signal-to-noise (SNR $\gtrsim$ 1000) [20]. In addition, unlike the compact binaries measured by LIGO, these sources will be in band for days to years, depending

on mass and redshift. This long time in band allows us to accurately track the wave phase and determine parameters of the MBH system with incredible accuracy. LISA parameter estimation of MBH binaries is the focus of Chapters 2 and 3.

Finally, supermassive black hole binaries ($M \sim 10^7 - 10^9 M_\odot$) lie in the very low frequency band. These GWs can be observed by studying the arrival times of pulsars [212, 67] and may form a confusion background [135]. The ultra low frequency band contains no binary sources. Instead, it consists of primordial gravitational waves amplified during inflation and imprinted on the cosmic microwave background (CMB). These waves can be detected with CMB polarization measurements [218, 137].[12]

1.4 Outline of this thesis

This thesis describes the observable effects of various general relativistic phenomena in compact binaries. In Chapter 2, the focus is on the coalescence of massive black hole binaries and their detection by LISA. The goal is to determine how well, given a detection of one of these events, we can extract the parameters of the system from the signal. Understanding the capabilities of the detector is critical for planning future research, as well as for properly promoting the mission while funding decisions are being made. We focus on a specific effect of general relativity: spin-induced precession. Post-Newtonian theory can be used to calculate the interactions of the black holes' spins with the orbital motion of the system and with each other. These interactions cause the spins and the orbital plane to precess, introducing various modulations into the waveform. In most previous analyses of LISA parameter estimation for MBHBs, these effects were left out; however, the modulations they generate can have a strong impact on the parameter measurement errors. In this chapter, based on [154] (see also [155] and [156]), we first define the problem in more detail. Then we present the form of the GWs emitted by the coalescing binaries, as well as the modulations

[12]Inflationary waves actually have a relatively flat spectrum [149], so they exist in all frequency bands; however, they are most detectable via CMB measurements. Stochastic waves from other sources, such as phase transitions, cosmic strings, and effects of extra dimensions, may also be found in multiple frequency bands; however, the existence of these waves is very uncertain.

added by the detector response. Next, we discuss how to estimate parameter errors using the Fisher matrix method. We then present comprehensive results for errors in mass, spin, sky position, and luminosity distance.

In Chapter 3, based on [157, 158], we continue our study into MBHB parameter estimation, focusing in more depth on LISA's ability to localize a source on the sky and in redshift. Recently there has been a great deal of interest in the potential to find electromagnetic counterparts to the MBHB GW events. Finding a counterpart could provide much more science output than a GW alone, especially if the counterpart is found during the coalescence, prior to merger. Therefore, we investigate the feasibility of *advanced* localization of the source, using only a portion of the inspiral GW signal. The chapter begins by discussing in more detail the importance of counterparts and some ideas about what form such counterparts might take. It then briefly reviews LISA's localization capability at merger, but in slightly more detail than in Chapter 2. Next, we describe a parallel study conducted by Kocsis et al. [144] which looks at the same problem of advanced localization, without precession effects but with a potentially troublesome approximation. We then present our own results for advanced localization. We conclude with a slightly different result—the dependence of our localization errors on the position of the source in the sky.

Chapter 4 changes gears to look at stellar-mass binaries which include neutron stars. We discuss briefly how measurements of a neutron star can help determine the equation of state of the dense matter at its center. One possible signature of the equation of state is the orbital radius at which a neutron star is tidally disrupted by a companion black hole. Traditional techniques for calculating this radius use Newtonian self-gravity for the stars, which is appropriate for white dwarfs but not for neutron stars. We propose a method for finding the structure of tidally distorted stars with relativistic self-gravity. Following work by Thorne and Campolattaro [239], we add nonradial perturbations to the standard Tolman-Oppenheimer-Volkoff metric for an unperturbed star. Then these perturbations are matched to a specific tidal field by analogy to work by Alvi [8]. We present some preliminary results for tidal disruption and comment on the effectiveness of our method. We then back away from

specific tidal fields and look instead at the general static response of a neutron star to a tidal field, which is characterized by a quantity known as the Love number. The gravitational wave phase of a compact binary is affected by the tidal excitation of oscillation modes, and to lowest order, this effect depends only on the Love number of the star. We calculate the Love numbers for various equations of state and then estimate the effect on the GW phase for neutron star-neutron star binaries.

Chapter 2

Measuring massive black hole binaries with gravitational waves

2.1 Introduction

2.1.1 Background to this analysis

Massive black hole (MBH) coalescences are one of the most exciting gravitational wave (GW) sources in the low-frequency band of the proposed Laser Interferometer Space Antenna (LISA). Merger tree calculations have shown that tens of events per year are likely to be detected [220, 175] with high SNR [20]. If these predictions hold true, merely *detecting* a MBH binary GW signal with LISA will prove routine. It will be far more interesting to use the information encoded in the GWs to learn about the sources. Some particularly important properties are the masses of the binary's members, their spins, the binary's location on the sky, and its distance from the solar system barycenter. Measuring a population of coalescence events could provide a wealth of data on the cosmological distribution and evolution of black hole masses and spins, while localization of a particular coalescence could lead to the identification of an electromagnetic counterpart. By providing specific information on source parameters, rather than just announcing the presence of a source, LISA will usher in the era of "gravitational wave astronomy." This is similar to the development

of neutrino astronomy, or even earlier, the expansion of astronomy into the radio and X-ray bands.

Using models for the MBH gravitational waveform and the detector response, we can calculate in advance how well the system parameters will be estimated by LISA. Such calculations are vital to members of the LISA team, who are trying to find the optimum design for the mission [231] and then secure funding for that design. In addition, scientists can use the results in planning how to best maximize the scientific output from LISA and contemporaneous astronomical instruments. Several studies of LISA's parameter estimation capability have been performed [62, 127, 29], differing from each other either in the particular parameter focus (e.g., masses or localization) or in the complexity of the signal and detector models. Recently, the problem has become so important that the LISA Science Team has established a new taskforce on LISA Parameter Estimation (LISA PE) [15].

These analyses have found that certain subsets of parameters tend to be highly correlated with each other, increasing the errors in parameter estimation. One such subset comprises the binary's "chirp mass" $\mathcal{M}$, its reduced mass μ, and the spin parameters β and σ (which are written out explicitly in Sec. 2.2.1). These four parameters influence the GW phase Φ. As discussed in [63, 194], the correlation coefficient between μ and β is nearly 1. It is thus difficult to "detangle" these parameters from one another in a measurement.

Another such subset consists of a binary's sky position, orientation, and luminosity distance D_L. To see why these parameters are strongly correlated, consider the form of the two polarizations of the strongest quadrupole harmonic of the gravitational waveform:

$$h_+(t) = 2\frac{\mathcal{M}^{5/3}(\pi f)^{2/3}}{D_L}(1+\cos^2\iota)\cos\Phi(t)\,, \tag{2.1}$$

$$h_\times(t) = -4\frac{\mathcal{M}^{5/3}(\pi f)^{2/3}}{D_L}\cos\iota\sin\Phi(t)\,. \tag{2.2}$$

The quantity ι is the binary's inclination relative to the line of sight: $\cos\iota \equiv \hat{\mathbf{L}}\cdot\hat{\mathbf{n}}$, where $\hat{\mathbf{L}}$, the direction of the binary's orbital angular momentum, defines its ori-

entation and $\hat{\mathbf{n}}$ is the direction from observer to source. The frequency $f(t) \equiv (1/2\pi)d\Phi/dt$.

One does not measure the polarizations h_+ and $h_\times$ directly; rather, one measures a sum $h_M(t)$ in which the two polarizations are weighted by antenna response functions F_+ and $F_\times$ as follows:

$$h_M(t) = F_+(\theta_N, \phi_N, \psi_N)h_+(t) + F_\times(\theta_N, \phi_N, \psi_N)h_\times(t) \ . \tag{2.3}$$

(This equation should be taken as schematic; see Sec. 2.2.3 for a more detailed and definitive description.) The angles θ_N and ϕ_N denote the location of the source on the sky in some appropriate coordinate system. The angle ψ_N, known as the "polarization angle," fixes the orientation of the component of $\hat{\mathbf{L}}$ perpendicular to the line of sight. (In other words, $\hat{\mathbf{L}}$ is fixed by ι and ψ_N.)

Measuring the phase determines chirp mass with high accuracy; the fractional error in $\mathcal{M}$ is often $\sim 10^{-3} - 10^{-4}$. As far as amplitude is concerned, the chirp mass can be regarded as measured exactly. What remains is to determine, from the measured amplitude and the known $\mathcal{M}$, the angles θ_N, ϕ_N, ψ_N, ι, and the distance D_L.

As (2.1), (2.2), and (2.3) illustrate, these five parameters are strongly correlated. The motion of LISA around the Sun breaks these degeneracies to some extent — the angles θ_N and ϕ_N appearing in (2.3) can be regarded as best defined in a coordinate system tied to LISA. As the antenna orbits the Sun, these angles become effectively time dependent. The one-year periodicity imposed by this motion makes it possible to detangle these parameters. Analyses typically find that the position of a coalescence event at $z \sim 1$ can be determined, on average, to an ellipse which is $1.5 - 2$ degrees across in the long direction and $1.5 - 2$ times smaller in the short direction[1] [62, 29, 123]. The distance to such a binary can be determined to $1\% - 2\%$ accuracy on average (less in some exceptional cases) [127, 29, 123].

[1]It is worth bearing in mind that the full moon subtends an angle of about 30 arcminutes, while the Hubble Deep Field [256] is only ~ 2.5 arcminutes on each side.

2.1.2 Black hole spin and spin precession

These previous analyses all ignore an important piece of relativistic physics: the precession of each binary member's spin vector due to its interaction with the spacetime in which it moves. In general relativity, the spacetime of an isolated object can be regarded as having an "electric piece," arising from the object's mass and mass distribution, and a "magnetic piece," arising from the object's mass currents and their distribution.[2] Spin precession consists of a "geodetic" term, arising from the parallel transport of the spin vector in the gravitoelectric field of the other hole, and "Lense-Thirring" terms, caused by the gravitomagnetic field of the other hole. The basic physics of gravitomagnetic precession can be simply understood by analogy with a similar (and closely related) electromagnetic phenomenon — the precession of a magnetic dipole $\boldsymbol{\mu}$ immersed in an external magnetic field $\mathbf{B}$. An object's spin angular momentum $\mathbf{S}$ can be regarded as a gravitational "magnetic dipole." When immersed in a "gravitomagnetic field," one finds that $\mathbf{S}$ feels a torque, just as a magnetic dipole $\boldsymbol{\mu}$ experiences a torque when immersed in magnetic field $\mathbf{B}$. In a binary black hole system, the gravitomagnetic field arises from the binary's orbital motion and the spins of its members. Precession thus includes both spin-orbit (geodetic and orbital gravitomagnetic) and spin-spin effects [241].

As the spins precess, they do so in such a way that the *total* angular momentum $\mathbf{J} = \mathbf{L} + \mathbf{S}_1 + \mathbf{S}_2$ is held constant; the orbital angular momentum $\mathbf{L}$ precesses to compensate for changes in $\mathbf{S}_1$ and $\mathbf{S}_2$. As a consequence, the inclination angle ι and polarization angle ψ_N become time varying (as do certain other quantities appearing in the GW phase function Φ). Figure 2-1 shows the so-called "polarization amplitude," defined in Sec. 2.2.3, of the waveform measured by a particular detector. Without precession, this quantity is modulated by the orbital motion of LISA, helping to provide some information about the binary's sky position. The polarization amplitude also depends on the angles ι and ψ_N, so it undergoes additional modu-

[2]This analogy is most apt in the weak field. In that limit, one can recast the Einstein field equations of general relativity into a form quite similar to Maxwell's equations; see [241] for detailed discussion. Though the analogy does not fit quite so well in strong-field regions, it remains accurate enough to be useful.

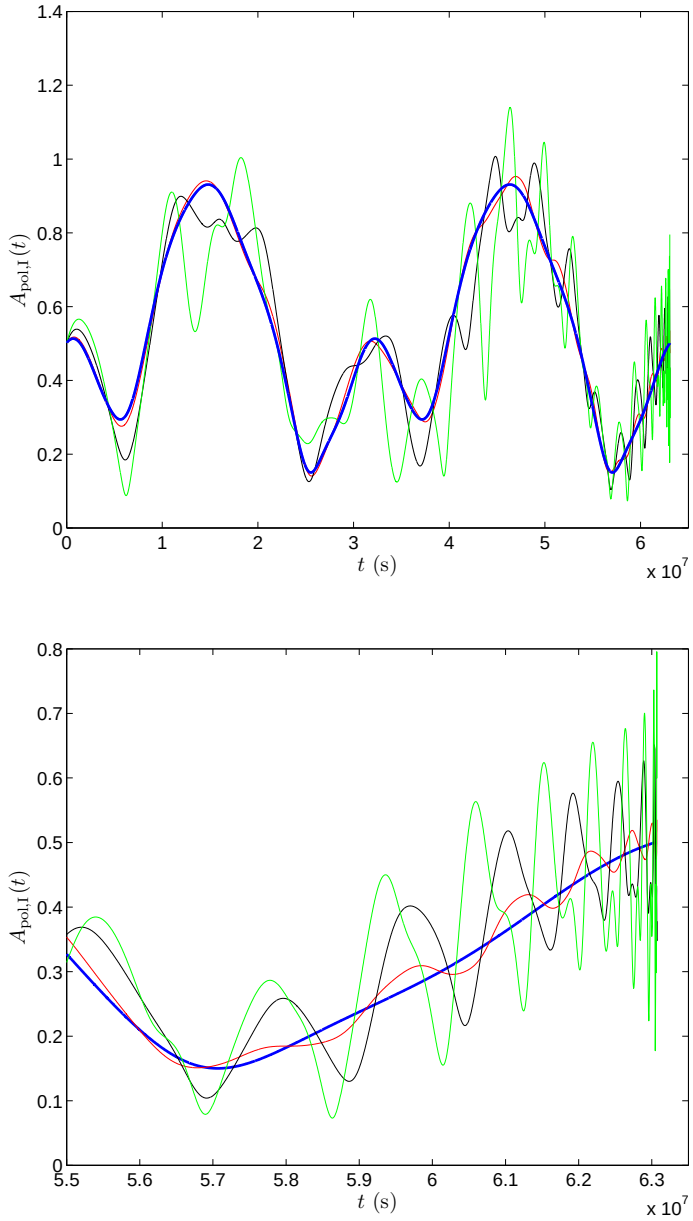

Figure 2-1: These figures depict the "polarization amplitude" $A_{\mathrm{pol}}(t)$ of a signal measured in detector I for a selection of spins: $\chi_1 = \chi_2 = 0$ (blue line), $\chi_1 = \chi_2 = 0.1$ (red line), $\chi_1 = \chi_2 = 0.5$ (black line), and $\chi_1 = \chi_2 = 0.9$ (green line). ($\chi = |\mathbf{S}|/m^2$ is the dimensionless spin parameter.) The top figure shows the final two years of inspiral. The spinless curve has periodicity of one year, corresponding to the motion of LISA around the Sun. Notice that as spin is introduced, the curves become more strongly modulated, with the number of additional oscillations growing as the spin is increased. By tracking these spin-precession-induced modulations, it becomes possible to better measure parameters like mass and sky position and measure spin for the first time. The bottom figure shows a close-up of the final months of inspiral. Precession effects increase drastically as merger approaches.

lation when precession is included. Such precession-imposed time variations quite thoroughly break many of the degeneracies which have been found to limit parameter measurement accuracy in earlier analyses.

It is without a doubt that black holes in nature spin. Observations are not yet precise enough to indicate the value of typical black hole spins; the evidence to date does, however, seem to indicate that fairly rapid rotation is common. For example, the existence of jets from active systems seems to require non-negligible black hole spin — jets appear to be "launched" by the shearing of magnetic field lines (supported by the highly conductive, ionized material accreting onto the black hole) by the differential rotation of spacetime around a rotating black hole [37, 174]. Also, observations of highly distorted iron K-α lines — a very sharp fluorescence feature in the rest frame of the emitting iron ions — indicate that this emission is coming from very deep within a gravitational potential (at radii less than that of the *Schwarzschild* innermost stable circular orbit $r = 6M$) and is smeared by near luminal relativistic speeds to boot [205]. Though perhaps influenced somewhat by selection effects[3], these pieces of evidence are strong hints that the black holes which will form the binaries we hope to measure will be strongly influenced by spin.

The only limit in which spin precession can be neglected is that in which the spins of the binary's members are exactly parallel (or antiparallel) to one another and to the orbital angular momentum $\mathbf{L}$. Since the target binaries of this analysis are created by galactic merger processes, their members will almost certainly have no preferred alignment — random spin and orbit orientation is expected to be the rule. (This expectation is borne out by work [213] showing that jets in active galaxies are oriented randomly with respect to the disks of their host galaxies.) Taking into account spin precession is thus of paramount importance for GW observations of merging black hole systems.

A great deal of work has gone into developing families of model waveforms ("templates") sufficiently robust to *detect* GWs from spinning and precessing binaries, at

[3]The systems for which we have constraints on spin are systems which are actively accreting and thus most likely to be rapidly spinning [131].

least in the context of measurements by ground-based detectors [43, 187, 42, 41, 106, 105, 104]. The key issue in this case is that the various modulations on the waveform imposed by the binary's precession smear its power over a wider spectral range, making it much more difficult to detect at the (relatively) low signal-to-noise ratios (SNRs) expected for ground-based observations. Not as much work has gone into the complementary problem of *measuring* these waves — examining the impact precession has upon the precision with which binary properties may be inferred from the waves. To date, the most complete and important analysis of this type is that of Vecchio [250]. Vecchio focuses (for simplicity) on equal-mass binaries and only includes the leading "spin-orbit" precession term. This limit is particularly nice as a first analysis of this problem, since it can be treated (largely) analytically (cf. discussion in Sec. IIIB of [250]).

Vecchio's work largely confirms the intuitive expectation discussed above — the precision with which masses are measured is substantially improved; in particular, the reduced mass of the system can be measured with several orders of magnitude more accuracy. Parameters such as the sky location of the binary and the luminosity distance are also measured more accurately, but only by a factor of 2 – 10.

2.1.3 This analysis

The goal here is to update Vecchio's pioneering analysis by taking the precession equations and the wave phase to the next higher order and by performing a broader parameter survey (including the impact of mass ratio). By taking the precession equations to higher order, we include "spin-spin" effects — precessional effects due to one black hole's spin interacting with gravitomagnetic fields from the other hole's spin. By taking the wave phase to higher order, we include, among other terms, a time-dependent spin-spin interaction. Finally, when the mass ratio differs from 1, the geodetic spin-orbit term causes the two spins to precess at different rates, even without the spin-spin corrections.

Including these effects means that the precession *cannot* be modeled with a simple, analytic rule — the equations of precession must be integrated numerically as inspiral

proceeds, incurring a significant performance cost. Fortunately, the basic "engine" on which this code is based [127] runs extremely fast, thanks largely to the use of spectral integrators (which, in turn, is thanks to a suggestion by E. Berti [29]), so total run time remains reasonable.

The cost in efficiency due to the inclusion of higher-order effects is offset by the more complete description of the signal they provide. An important consequence is that it now becomes possible from GW measurements to determine the spin of each member of the binary. With Vecchio's approximations, only three components of the black holes' vector spins can be determined — enough to constrain, but not determine, their spin magnitudes. The more general approach of this chapter allows for the measurement of all six vector spin components. It is, therefore, the first analysis indicating how well spin magnitude can be measured from merging comparable-mass binary systems. (As Barack and Cutler have shown [21], spin is very well determined by measurements of GWs from *extreme mass ratio* binaries.)

The error estimates are computed using the maximum likelihood formalism first introduced in the context of GW measurements by Finn [90]. A potential worry is that we are using a Gaussian approximation to the likelihood function. This approximation is very convenient since it allows us to directly compute a Fisher information matrix. Its inverse is the covariance matrix, which directly encodes the estimated 1-σ errors in measured parameters, as well as correlations among different parameters. The Gaussian approximation is known to be accurate when the SNR is "high enough" [90, 63].

Unfortunately, it is not particularly obvious what "high enough" really means. In our case, we are estimating measurement errors on 15 parameters[4] — a rather fearsome number to fit. The Gaussian approximation almost certainly *underestimates* measurement error, since it assumes the likelihood function is completely determined by its curvature in the vicinity of a maximum, missing the possibility of a long tail to large error. We thus fear that our estimates are likely to be optimistic, especially

[4] 2 masses; 2 angles specifying the initial orientation of the binary's orbit; 4 angles specifying the initial orientation of the spins; 2 spin magnitudes; the time at which coalescence occurs; the phase at coalescence; 2 angles specifying the binary's position on the sky; and the distance to the binary.

for events with relatively small SNR. It would be a good idea to directly compute the likelihood function in a few important corners of parameter space and compare to the Gaussian predictions. This would both quantify the degree to which our calculations are too optimistic and help to determine how large SNR must be for this approximation to be reliable.

In addition to concerns about the Gaussian approximation, it must be noted that the waveform family we use for our analysis is somewhat limited. We use a post-Newtonian description of the GWs from these binaries. Since our analysis requires us to follow these binaries deep into the strong field where the usual post-Newtonian expansion is likely to be somewhat unreliable, it is likely that we are introducing some systematic error. In more recent work, this problem has been avoided by stitching the post-Newtonian waveform to a waveform generated by numerical relativity when the latter becomes more accurate [242].

Even within the post-Newtonian description, we have made various approximations. To begin with, we use the so-called "restricted post-Newtonian" waveform. In this approximation, we compute the phase to a desired post-Newtonian order but keep only the lowest order amplitude term, the Newtonian quadrupole of (1.37). It has been recognized for some time that additional information is carried by higher order terms [113]. Recent work [16, 245, 196] has confirmed that parameters are more sharply constrained when these terms are included in the wave model.

We also only calculate the phase to second post-Newtonian (2PN) order, less than the current standard of 3.5PN [34]. In addition, the equations of spin precession that we use are only given to the leading order needed to see spin-orbit and spin-spin precession effects [10]. Higher spin-orbit corrections to the equations of motion and precession have recently been derived [86], as have their impact on the the waves' phasing [35]. Another analysis [197] has worked out higher-order spin-spin corrections to the post-Newtonian metric, from which it would not be too difficult to work out equations of motion and precession and then the modification to the waves' phase.

Finally, it should be noted that the frequency domain expression of the signal which we use is derived formally using a "stationary phase" approximation, which

in general is known to be good for nonprecessing binaries [72]. This approximation is based on the idea that the binary's orbital frequency is changing "slowly." The orbital frequency is thus well-defined over "short" time scales. Quantitatively, this amounts to a requirement that the time scale on which radiation reaction changes the orbital frequency, $T_{\rm insp}$, be much longer than an orbital period, $T_{\rm orb}$. Precession introduces a new time scale, $T_{\rm prec}$, the time it takes for the angular momentum vectors to significantly change their orientations. For the stationary phase approximation to be accurate when precession physics is included in the waveform, we must in addition require $T_{\rm prec} \gg T_{\rm orb}$, a somewhat more stringent requirement than $T_{\rm insp} \gg T_{\rm orb}$. No doubt, a certain amount of error is introduced due to the breakdown of this condition late in the inspiral.

Thus, the results which we present here should be taken as *indicative* of how well LISA is likely to be able to measure the parameters of massive black hole binaries, but cannot be considered definitive. We are confident however that the *improvement* in measurement accuracy obtained by taking spin precession into account is robust. Specifically, we see that errors in masses are reduced dramatically, from one to several orders of magnitude. Errors in sky position and distance are also reduced, but by a smaller factor. Such improvement may nonetheless critically improve the ability of LISA to interface with electromagnetic observatories [142, 144, 157, 143]. Finally, the added information in the precession signal allows us to measure the spins of the holes. These improvements due to precession will certainly survive and play an important role even in an analysis which addresses the caveats we list above.

2.1.4 Outline of this chapter

The remainder of the chapter is organized as follows. In Sec. 2.2, we discuss the gravitational waveform generated by binary black hole coalescence, focusing on the slow, adiabatic inspiral. Section 2.2.1 describes the "intrinsic" waveform produced by the motion of the orbiting black holes as given in the restricted post-Newtonian expansion of general relativity. Section 2.2.2 then describes the post-Newtonian precession equations which we use to model the evolution of the spins of a binary's

members, as well as how those precessions influence the waveform. Finally, in Sec. 2.2.3 we describe "extrinsic" effects which enter the measured waveform through its measurement by the LISA constellation.

In Sec. 2.3 we describe our parameter estimation formalism. Section 2.3.1 first summarizes the Fisher matrix method we use to estimate measurement errors. In Sec. 2.3.2, we then describe our model for the noise which we expect to accompany LISA measurements.

Section 2.4 presents our results. After describing some critical procedural issues in the setup of our calculations in Sec. 2.4.1, we summarize our results for parameters intrinsic to the binary (particularly masses and spins) in Sec. 2.4.2 and for extrinsic parameters (particularly sky position and luminosity distance) in Sec. 2.4.3. In both cases, we compare, when appropriate, to results from a code which does not incorporate spin-precession physics. (This code was originally developed for the analysis presented in [127].) The general rule of thumb we find is that the accuracy with which masses can be determined is improved by about one to several orders of magnitude when precession physics is taken into account. In addition, we find that for low redshift ($z \sim 1$) binaries LISA should be able to determine the spins of the constituent black holes with a relative precision of $0.1\% - 10\%$, depending (rather strongly) on the spin value. Likewise, we find improvement in the measurement accuracy of extrinsic parameters, though not quite as striking — half an order of magnitude improvement in source localization and distance determination is a good, rough rule of thumb.

An important consequence of these improvements is that LISA should be able to localize low-redshift binaries — using GW measurements alone — to an ellipse that is perhaps a few $\times$ 10 arcminutes across in its widest direction and a factor of $2 - 4$ smaller along its minor axis. For higher redshift binaries ($z \sim 3 - 5$), this ellipse is several times larger, perhaps a few degrees in the long direction and tens of arcminutes to a degree or two in the narrow one. These results suggest that it should not be too arduous a task to search for electromagnetic counterparts to a coalescing binary black hole's GW signal [142, 144, 157, 143] — particularly at low redshift, these error ellipse sizes are comparable to the field of view of planned large-scale surveys.

A concluding and summarizing discussion is given in Sec. 2.5. Along with summarizing our major results and findings, we discuss future work which could allow us to quantitatively assess the consequences of some of the simplifying assumptions we have made.

At several points in this analysis, we need to convert between a source's redshift z and luminosity distance D_L. To make this conversion, we assume a flat cosmology ($\Omega_{\rm total} = 1$) with contributions from matter ($\Omega_M = 0.25$) and from a cosmological constant (equation of state parameter $w = -1$, $\Omega_\Lambda = 0.75$). We also choose a Hubble constant $H_0 = 75\,{\rm km\ s^{-1}\ Mpc^{-1}}$. These choices are in concordance with the latest fits presented by the WMAP team in their three-year analysis of the cosmic microwave background [229]. The luminosity distance as a function of redshift is then given by

$$D_L(z) = \frac{(1+z)c}{H_0} \int_0^z \frac{dz'}{\sqrt{\Omega_M (1+z')^3 + \Omega_\Lambda}}\,. \tag{2.4}$$

2.2 Gravitational waves from binary black hole inspiral

The GWs generated by a coalescing binary black hole system can be divided into three more or less distinct epochs [96]: (1) a slowly evolving *inspiral*, in which the black holes gradually spiral toward each other as the orbit decays due to GW emission; (2) a loud *merger*, in which the black holes come together and form a single body; and (3) a *ringdown*, in which the merged remnant of the binary settles down to its final state. Our analysis focuses on the inspiral, the most long-lived epoch of coalescence and the epoch in which spin precession plays a major dynamical role. Ringdown waves have been analyzed in other work [73, 90, 127, 71]; the most comprehensive recent analysis was performed by Berti, Cardoso, and Will [30]. The merger waveform, describing the strong-field and (potentially) violent process of the two black holes merging into a single body, has historically been poorly understood. Recent breakthroughs in numerical relativity have corrected this problem [201, 45, 19], and merger waveforms are now being used for parameter estimation [18, 242].

The inspiral waveform which will be measured by LISA is a combination of the intrinsic waveform created by the source and extrinsic features related to its location on the sky and modulation effects caused by the motion of the detector. In this section we review the relevant physics involved in the construction of the waveform.

For sources at cosmological distances, all time scales redshift by a factor $1+z$. In the $G = c = 1$ units that we use, all factors of mass enter as time scales; thus, masses are redshifted by this $1 + z$ factor. (Likewise, quantities such as spin which have dimension (time)^2 acquire a factor $(1 + z)^2$, etc.) In the equations written below, we do not explicitly write out these redshift factors; they should be taken to be implicit in all our equations. When discussing results, we will always quote masses as they would be measured in the rest frame of the source, with redshift given separately.

2.2.1 Intrinsic waveform

We treat the members of our binary as moving on quasi-circular orbits. Eccentricity is very rapidly bled away by gravitational radiation reaction [191], so it has traditionally been expected that these binaries will have essentially zero eccentricity by the time they enter LISA's frequency band (at least at the mass ratios we consider in this chapter, $1 \leq m_1/m_2 \leq 10$). However, dynamical scenarios that allow the binary to overcome the last parsec problem may also leave it with a significant eccentricity [12, 28]; it will be useful to include this effect in future studies.

We use the post-Newtonian formalism, an expansion in internal gravitational potential $U \sim M/r$ and internal source velocity $v \sim \sqrt{M/r}$, to build our waveforms. Post-Newtonian orders are often defined in terms of powers of v; corrections at nth post-Newtonian order go like v^{2n}. A detailed review of the post-Newtonian formalism can be found in the article by Blanchet [34]; the key pieces which we will use can be found in [34, 140, 36, 139, 255].

The post-Newtonian equations of motion, taken to second post-Newtonian (2PN) order, yield the following generalization of Kepler's third law relating orbital angular

frequency Ω and orbital radius[5] r [36]:

$$\Omega^2 = \frac{M}{r^3}\left[1-(3-\eta)\left(\frac{M}{r}\right)-\sum_{i=1}^{2}\left(2\frac{m_i^2}{M^2}+3\eta\right)\frac{\hat{\mathbf{L}}\cdot\mathbf{S}_i}{m_i^2}\left(\frac{M}{r}\right)^{3/2} + \left(6+\frac{41}{4}\eta+\eta^2-\frac{3}{2}\frac{\eta}{m_1^2m_2^2}[\mathbf{S}_1\cdot\mathbf{S}_2-3(\hat{\mathbf{L}}\cdot\mathbf{S}_1)(\hat{\mathbf{L}}\cdot\mathbf{S}_2)]\right)\left(\frac{M}{r}\right)^2\right]. \tag{2.5}$$

Here $M = m_1+m_2$ is the total mass of the system, and $\eta = \mu/M$, where $\mu = m_1m_2/M$ is the reduced mass. $\hat{\mathbf{L}}$ is the direction of the orbital angular momentum, and $\mathbf{S}_i$ is the spin angular momentum of black hole i. The magnitude of the spin can be expressed as $S_i = \chi_i m_i^2$, where $0 \le \chi_i \le 1$. The leading term is the standard result from Newtonian gravity. The $O(M/r)$ term is the first post-Newtonian correction; this is the same physics that, in solar system dynamics, causes the precession of the perihelion of Mercury. The $O((M/r)^{3/2})$ term contains spin-orbit corrections to the equation of motion. Finally, the $O((M/r)^2)$ term is a 2PN correction, which also includes spin-spin terms. From the equations of motion, the orbital energy of the binary E can also be computed [36]:

$$E = -\frac{\mu M}{2r}\left[1-\frac{1}{4}(7-\eta)\frac{M}{r}+\sum_{i=1}^{2}\left(2\frac{m_i^2}{M^2}+\eta\right)\frac{\hat{\mathbf{L}}\cdot\mathbf{S}_i}{m_i^2}\left(\frac{M}{r}\right)^{3/2} + \left(-\frac{7}{8}+\frac{49}{8}\eta+\frac{1}{8}\eta^2+\frac{\eta}{2m_1^2m_2^2}[\mathbf{S}_1\cdot\mathbf{S}_2-3(\hat{\mathbf{L}}\cdot\mathbf{S}_1)(\hat{\mathbf{L}}\cdot\mathbf{S}_2)]\right)\left(\frac{M}{r}\right)^2\right]. \tag{2.6}$$

The binary loses energy to gravitational waves at the rate [36]

$$\frac{dE}{dt} = -\frac{32}{5}\eta^2\left(\frac{M}{r}\right)^5\left[1-\left(\frac{2927}{336}+\frac{5}{4}\eta\right)\frac{M}{r} + \left(4\pi-\frac{1}{12}\sum_{i=1}^{2}\left[73\frac{m_i^2}{M^2}+75\eta\right]\frac{\hat{\mathbf{L}}\cdot\mathbf{S}_i}{m_i^2}\right)\left(\frac{M}{r}\right)^{3/2} + \left(\frac{293383}{9072}+\frac{380}{9}\eta-\frac{\eta}{48m_1^2m_2^2}[223\mathbf{S}_1\cdot\mathbf{S}_2-649(\hat{\mathbf{L}}\cdot\mathbf{S}_1)(\hat{\mathbf{L}}\cdot\mathbf{S}_2)]\right)\left(\frac{M}{r}\right)^2\right]. \tag{2.7}$$

[5]The radius is given in harmonic coordinates, which satisfy $\Box x^\alpha = 0$ when treated as four scalar fields. Here, unlike in the linearized gravity discussion of Chapter 1, $\Box \equiv g^{\mu\nu}\partial_\mu\partial_\nu$.)

As mentioned above, we actually use the "restricted" 2PN waveform. This approximation can be understood by writing the waveform (somewhat schematically) as [63]

$$h(t) = \mathrm{Re}\left(\sum_{x,m} h_m^x(t) e^{im\Phi_{\mathrm{orb}}(t)}\right), \tag{2.8}$$

where x labels PN order, m is a harmonic index, and $\Phi_{\mathrm{orb}}(t) = \int^t \Omega(t')dt'$ is orbital phase. In the restricted post-Newtonian waveform, we throw out all amplitude terms except h_2^0 (the "Newtonian quadrupole" term) but compute $\Phi_{\mathrm{orb}}(t)$ to some specified PN order. The restricted PN approximation is motivated by the fact that matched filtering — matching a signal in noisy data by cross-correlating with a theoretical template — is much more sensitive to phase information than to the amplitude. Since the h_2^0 harmonic contributes most strongly to the waveform over most of the inspiral, the restricted PN approximation is expected to capture the most important portion of the inspiral waveform. However, as mentioned earlier, the subleading terms do carry important information. We plan to include these terms in future work, examining how they combine with precession to determine system parameters.

At any rate, within the restricted PN approximation, the waveform can be written

$$h_{ij}(t, \mathbf{x}) = -\frac{4\mathcal{M}^{5/3}(\pi f)^{2/3}}{|\mathbf{x}|} \begin{bmatrix} \cos\Phi(t) & \sin\Phi(t) & 0 \\ \sin\Phi(t) & -\cos\Phi(t) & 0 \\ 0 & 0 & 0 \end{bmatrix}, \tag{2.9}$$

where $|\mathbf{x}|$ is the distance to the binary, $\mathcal{M} = \mu^{3/5}M^{2/5}$ is the "chirp mass" (so called because it largely determines the rate at which the system's frequency evolves, or "chirps"), $f = \Omega/\pi = 2f_{\mathrm{orb}}$ is the GW frequency, and $\Phi(t) = \int^t 2\pi f(t')dt' = 2\Phi_{\mathrm{orb}}$ is the GW phase. We have chosen a coordinate system oriented such that the binary's orbit lies within the xy-plane; this tensor will later be projected onto polarization basis tensors to construct the measured polarizations h_+ and $h_\times$.

The rate at which the frequency changes due to the emission of gravitational

radiation can be found using (2.5)-(2.7) (with $f = \Omega/\pi$):

$$\frac{df}{dt} = \frac{96}{5\pi\mathcal{M}^2}(\pi\mathcal{M}f)^{11/3}\left[1 - \left(\frac{743}{336} + \frac{11}{4}\eta\right)(\pi Mf)^{2/3} + (4\pi - \beta)(\pi Mf) + \left(\frac{34103}{18144} + \frac{13661}{2016}\eta + \frac{59}{18}\eta^2 + \sigma\right)(\pi Mf)^{4/3}\right]. \tag{2.10}$$

Notice that the chirp mass $\mathcal{M}$ dominates the rate of change of f; the reduced mass μ and parameters β and σ have an influence as well. The parameter β describes spin-orbit interactions and is given by

$$\beta = \frac{1}{12}\sum_{i=1}^{2}\left[113\left(\frac{m_i}{M}\right)^2 + 75\frac{\mu}{M}\right]\frac{\hat{\mathbf{L}}\cdot\mathbf{S}_i}{m_i^2}. \tag{2.11}$$

The parameter σ describes spin-spin interactions:

$$\sigma = \frac{\mu}{48M(m_1^2m_2^2)}[721(\hat{\mathbf{L}}\cdot\mathbf{S}_1)(\hat{\mathbf{L}}\cdot\mathbf{S}_2) - 247(\mathbf{S}_1\cdot\mathbf{S}_2)]. \tag{2.12}$$

Notice that β and σ depend on the angles between the binary's angular momentum and the two spins. In previous analyses which have neglected precession, β and σ are constants; precession makes them time dependent.

Using (2.10), we can now integrate to find[6]

$$t(f) = t_c - \frac{5}{256}\mathcal{M}(\pi\mathcal{M}f)^{-8/3}\left[1 + \frac{4}{3}\left(\frac{743}{336} + \frac{11}{4}\eta\right)(\pi Mf)^{2/3} - \frac{8}{5}(4\pi - \beta)(\pi Mf) + 2\left(\frac{3058673}{1016064} + \frac{5429}{1008}\eta + \frac{617}{144}\eta^2 - \sigma\right)(\pi Mf)^{4/3}\right]. \tag{2.13}$$

The parameter t_c formally defines the time at which f diverges within the post-Newtonian framework. In reality, we expect finite-size effects to significantly modify the binary's evolution as the members come into contact. The system evolves so

[6]This expression and the next are derived by assuming that the spins are constant. We then plug the time-dependent spins into the final results. Instead, we should have put the time-dependent spins in at a lower level (i.e., (2.10)) and numerically integrated to find the time and phase. This method is now being used in follow-up work which uses some of our code [15].

quickly as the bodies come together that t_c is nonetheless a useful surrogate for a "time of coalescence". Finally, the wave phase $\Phi(t) = \int^t 2\pi f(t')dt'$ as a function of f is given by

$$\Phi(f) \equiv \Phi[t(f)] = \Phi_c - \frac{1}{16}(\pi\mathcal{M}f)^{-5/3}\left[1 + \frac{5}{3}\left(\frac{743}{336} + \frac{11}{4}\eta\right)(\pi Mf)^{2/3} \right.$$
$$\left. -\frac{5}{2}(4\pi - \beta)(\pi Mf) + 5\left(\frac{3058673}{1016064} + \frac{5429}{1008}\eta + \frac{617}{144}\eta^2 - \sigma\right)(\pi Mf)^{4/3}\right], \tag{2.14}$$

where Φ_c is the phase at time t_c. The restricted PN waveform is then constructed by inserting (2.14) into (2.9).

2.2.2 Precession equations

We next examine the effects of precession on the binary system. As discussed in Sec. 2.1, spin-orbit and spin-spin interactions cause the black hole spins $\mathbf{S}_1$ and $\mathbf{S}_2$ to precess. Precession occurs, at leading order, on a time scale $T_{\rm prec} \propto r^{5/2}$ at large separations [214]. Since this is smaller than the inspiral time scale $T_{\rm insp} \propto r^4$, we treat the total angular momentum $\mathbf{J} = \mathbf{L} + \mathbf{S}_1 + \mathbf{S}_2$ as constant over $T_{\rm prec}$. The orbital angular momentum $\mathbf{L}$ must then precess to compensate for changes in $\mathbf{S}_1$ and $\mathbf{S}_2$. Since $T_{\rm prec}$ is longer than the orbital time scale $T_{\rm orb} \propto r^{3/2}$, we use an orbit-averaged version of the precession equations[7] [10, 139]:

$$\dot{\mathbf{S}}_1 = \frac{1}{r^3}\left[\left(2 + \frac{3}{2}\frac{m_2}{m_1}\right)\mu\sqrt{Mr}\hat{\mathbf{L}}\right] \times \mathbf{S}_1 + \frac{1}{r^3}\left[\frac{1}{2}\mathbf{S}_2 - \frac{3}{2}(\mathbf{S}_2 \cdot \hat{\mathbf{L}})\hat{\mathbf{L}}\right] \times \mathbf{S}_1\,, \tag{2.15}$$

$$\dot{\mathbf{S}}_2 = \frac{1}{r^3}\left[\left(2 + \frac{3}{2}\frac{m_1}{m_2}\right)\mu\sqrt{Mr}\hat{\mathbf{L}}\right] \times \mathbf{S}_2 + \frac{1}{r^3}\left[\frac{1}{2}\mathbf{S}_1 - \frac{3}{2}(\mathbf{S}_1 \cdot \hat{\mathbf{L}})\hat{\mathbf{L}}\right] \times \mathbf{S}_2\,, \tag{2.16}$$

where dots over quantities denote time derivatives and[8] $r = M^{1/3}/(\pi f)^{2/3}$. These equations each have two pieces [241]. Consider the equation for $\dot{\mathbf{S}}_1$. The first piece, which contains no $\mathbf{S}_2$ dependence, is the spin-orbit term. This term, which comes in

[7] In fact, orbit-averaging is necessary for the existence of quasi-circular orbits in the two-spin case [139].

[8] We use only the lowest-order Newtonian orbital separation in these equations. Including more terms would introduce higher-order effects into the precession.

at 1PN order, is due to the geodetic precession of $\mathbf{S}_1$ as hole 1 orbits in the spacetime generated by the mass of hole 2, and to the Lense-Thirring precession of $\mathbf{S}_1$ in the gravitomagnetic field generated by the orbital motion of hole 2. The second piece is the spin-spin term, which enters at 1.5PN order. This term can be understood as the Lense-Thirring precession of $\mathbf{S}_1$ in the gravitomagnetic field generated by the spin of hole 2. Note that the magnitudes of the spins do not change at this order; see [10] for more details. From conservation of total angular momentum on short time scales, we have

$$\dot{\mathbf{L}} = -(\dot{\mathbf{S}}_1 + \dot{\mathbf{S}}_2)\,. \tag{2.17}$$

Over longer time scales, we must also consider the change in total angular momentum due to the radiation reaction, which is given by

$$\dot{\mathbf{J}} = -\frac{32}{5}\frac{\mu^2}{r}\left(\frac{M}{r}\right)^{5/2}\hat{\mathbf{L}} \tag{2.18}$$

to lowest order.

Considering only the spin-orbit terms and taking the limit $S_2 = 0$ or $m_1 = m_2$ leads to a system whose precession can described analytically; this is the "simple precession" limit described in [10]. Simple precession can be visualized as a rotation of $\mathbf{L}$ and $\mathbf{S} = \mathbf{S}_1 + \mathbf{S}_2$ around the total angular momentum $\mathbf{J}$. (Since inspiral shrinks $\mathbf{J}$, the precession is actually around a slightly different direction $\mathbf{J}_0$; see [10] for further discussion.)

Since Vecchio restricts his analysis to $m_1 = m_2$ and does not include the spin-spin interaction, this limit is appropriate for his work [250]. As a consequence, Vecchio takes the quantities $|\mathbf{S}|$, $\hat{\mathbf{L}}\cdot\hat{\mathbf{S}}$, $\hat{\mathbf{S}}_1\cdot\hat{\mathbf{S}}_2$, and β to be constant. (He does not include the spin-spin term σ in the analysis.) Here, we will study the impact of the full (albeit orbit-averaged) precession equations, including spin-spin terms, and include the impact of mass ratio. An analytic description is not possible in this case, so we must integrate these equations numerically. The behavior is qualitatively similar to the simple precession case, but with significant quantitative differences. For example, β now oscillates around an average value. For unequal masses (say $m_1/m_2 \gtrsim 2$), the

difference due to precession can be substantial [139]. Such cases are also astrophysically the most interesting — a mass ratio of roughly 10 is favored in binary black hole formation scenarios arising from hierarchical structure formation [219].

At this point, we note that precession's effect on the waveform is to modify the functions β and σ which appear in the post-Newtonian phase (2.14) and time-frequency relation (2.13). In the next section, we consider extrinsic effects on the waveform and find that precession of the orbital plane modifies them as well.

2.2.3 Extrinsic effects

We have now constructed the intrinsic GWs emitted by a precessing binary in the restricted post-Newtonian approximation. The waveform measured by LISA will also include extrinsic effects due to the binary's location on the sky and the motion of the detector.

We can write the wave as a combination of two orthogonal polarizations propagating in the $-\hat{\mathbf{n}}$ direction (so that $\hat{\mathbf{n}}$ is the position of the binary on the sky). Define $\hat{\mathbf{p}}$ and $\hat{\mathbf{q}}$ as axes orthogonal to $\hat{\mathbf{n}}$, with $\hat{\mathbf{p}} = \hat{\mathbf{n}} \times \hat{\mathbf{L}}/|\hat{\mathbf{n}} \times \hat{\mathbf{L}}|$ and $\hat{\mathbf{q}} = \hat{\mathbf{p}} \times \hat{\mathbf{n}}$. These are the principal axes for the wave; that is, they are defined so that the two polarizations are exactly 90° out of phase. The polarization basis tensors for these axes are $H^+_{ij} = p_i p_j - q_i q_j$ and $H^\times_{ij} = p_i q_j + q_i p_j$:

$$h_{ij}(t) = h_+(t) H^+_{ij} + h_\times(t) H^\times_{ij} \,, \tag{2.19}$$

where

$$h_+(t) = 2\frac{\mathcal{M}^{5/3}(\pi f)^{2/3}}{D_L}[1 + (\hat{\mathbf{L}} \cdot \hat{\mathbf{n}})^2] \cos[\Phi(t) + \delta_p\Phi(t)] \,, \tag{2.20}$$

$$h_\times(t) = -4\frac{\mathcal{M}^{5/3}(\pi f)^{2/3}}{D_L}(\hat{\mathbf{L}} \cdot \hat{\mathbf{n}}) \sin[\Phi(t) + \delta_p\Phi(t)] \,. \tag{2.21}$$

Here D_L is the luminosity distance to the source. Notice that the weighting of the two polarizations depends upon the direction of the orbital angular momentum vector relative to the sky position. This weighting is time dependent when the orbital plane

precesses. The precession also causes an $\hat{\mathbf{n}}$-dependent change in the observed orbital phase $\Phi_{\rm orb}(t)$ [10]. Multiplying by a factor of 2, the change in the wave phase is

$$\delta_p \Phi(t) = -\int_t^{t_c} \delta_p \dot{\Phi}(t') dt' \,, \tag{2.22}$$

where

$$\delta_p \dot{\Phi}(t) = \frac{2\hat{\mathbf{L}} \cdot \hat{\mathbf{n}}}{1 - (\hat{\mathbf{L}} \cdot \hat{\mathbf{n}})^2} (\hat{\mathbf{L}} \times \hat{\mathbf{n}}) \cdot \dot{\hat{\mathbf{L}}} \,. \tag{2.23}$$

We now consider the GW as measured by the detector. All of this analysis is done using the long wavelength ($\lambda \gg L$, where L is the LISA arm length) approximation introduced by Cutler [62]; more details can be found there. This approximation is appropriate for our purposes since most of the signal accumulates at low frequencies where the wavelength is in fact greater than the arm length. The full LISA response function, including arm-length effects, is discussed in [61, 210].

LISA consists of three spacecraft arranged in an equilateral triangle, 5×10^6 km apart. The center of mass of the configuration orbits the Sun 20° behind the Earth. The triangle is oriented at 60° to the ecliptic, so the orbits of the individual spacecraft will all be in different planes. This causes the triangle to spin around itself as it orbits the Sun. Following Cutler, we define a barred "barycenter" coordinate system $(\bar{x}, \bar{y}, \bar{z})$, which is fixed in space with the $\bar{x}\bar{y}$-plane aligned with the ecliptic, and an unbarred "detector" coordinate system (x, y, z), which is attached to the detector. The z axis always points toward the Sun, 60° away from vertical, while the x and y axes pinwheel around it. A particular binary will have fixed coordinates in the barycenter system, but its detector coordinates will be time varying.

The three arms act as a pair of two-arm detectors. We are first interested in the strain measured in detector I, that formed by arms 1 and 2:

$$h_{\rm I}(t) = \frac{\delta L_1(t) - \delta L_2(t)}{L} \,, \tag{2.24}$$

where $\delta L_1(t)$ and $\delta L_2(t)$ are the differences in length in arms 1 and 2 as the wave passes. L is the unperturbed length of the arms. Using the geometry of the detector

and the equation of geodesic deviation [97], we find

$$h_{\rm I}(t) = \frac{\sqrt{3}}{2}\left[\frac{1}{2}(h_{xx} - h_{yy})\right] . \tag{2.25}$$

To obtain h_{xx} and h_{yy} for use in these equations, we must rotate the waveform from the principal axes into the detector frame. The result is that detector I measures both polarizations, modulated by the antenna pattern of that detector:

$$\begin{aligned} h_{\rm I}(t) = \frac{\sqrt{3}}{2}\frac{\mathcal{M}^{5/3}(\pi f)^{2/3}}{D_L}&(2[1 + (\hat{\mathbf{L}}\cdot\hat{\mathbf{n}})^2]F_{\rm I}^{+}(\theta_N,\phi_N,\psi_N)\cos[\Phi(t) + \delta_p\Phi(t)] \\ &- 4(\hat{\mathbf{L}}\cdot\hat{\mathbf{n}})F_{\rm I}^{\times}(\theta_N,\phi_N,\psi_N)\sin[\Phi(t) + \delta_p\Phi(t)]) . \end{aligned} \tag{2.26}$$

Detector I acts like a "standard" 90° GW interferometer (e.g. LIGO), with the response scaled by $\sqrt{3}/2$ (due to the 60° opening angle of the constellation). The antenna pattern functions are given by

$$F_{\rm I}^{+}(\theta_N,\phi_N,\psi_N) = \frac{1}{2}(1+\cos^2\theta_N)\cos 2\phi_N\cos 2\psi_N - \cos\theta_N\sin 2\phi_N\sin 2\psi_N , \tag{2.27}$$

$$F_{\rm I}^{\times}(\theta_N,\phi_N,\psi_N) = \frac{1}{2}(1+\cos^2\theta_N)\cos 2\phi_N\sin 2\psi_N + \cos\theta_N\sin 2\phi_N\cos 2\psi_N . \tag{2.28}$$

Here $\theta_N = \theta_N(t)$ and $\phi_N = \phi_N(t)$ are the spherical angles for the binary's direction in the detector frame. They are related to the *constant* barycenter angles $\bar{\theta}_N$ and $\bar{\phi}_N$ by the expressions

$$\cos\theta_N(t) = \frac{1}{2}\cos\bar{\theta}_N - \frac{\sqrt{3}}{2}\sin\bar{\theta}_N\cos(\bar{\Phi}_D(t) - \bar{\phi}_N) , \tag{2.29}$$

$$\phi_N(t) = \bar{\Phi}_D(t) - \frac{\pi}{2} + \tan^{-1}\left(\frac{\sqrt{3}\cos\bar{\theta}_N + \sin\bar{\theta}_N\cos[\bar{\Phi}_D(t) - \bar{\phi}_N]}{2\sin\bar{\theta}_N\sin[\bar{\Phi}_D(t) - \bar{\phi}_N]}\right) , \tag{2.30}$$

where $\bar{\Phi}_D(t) = 2\pi t/T_D \equiv 2\pi t/(1\text{ yr})$ is the orbital phase of the detector. (A comparison to Cutler shows that we have chosen $\bar{\Phi}_D(0) = 0$ and orientation parameter $\alpha_0 = -\pi/2$ as initial conditions.) $\psi_N = \psi_N(t)$ is the polarization angle of the wave

in the detector frame:

$$\tan\psi_N = \frac{\hat{\mathbf{q}}\cdot\hat{\mathbf{z}}}{\hat{\mathbf{p}}\cdot\hat{\mathbf{z}}} = \frac{\hat{\mathbf{L}}\cdot\hat{\mathbf{z}} - (\hat{\mathbf{L}}\cdot\hat{\mathbf{n}})(\hat{\mathbf{z}}\cdot\hat{\mathbf{n}})}{\hat{\mathbf{n}}\cdot(\hat{\mathbf{L}}\times\hat{\mathbf{z}})}\,, \tag{2.31}$$

where

$$\hat{\mathbf{L}}\cdot\hat{\mathbf{z}} = \frac{1}{2}\cos\bar{\theta}_L(t) - \frac{\sqrt{3}}{2}\sin\bar{\theta}_L(t)\cos(\bar{\Phi}_D(t) - \bar{\phi}_L(t))\,, \tag{2.32}$$

$$\hat{\mathbf{L}}\cdot\hat{\mathbf{n}} = \cos\bar{\theta}_L(t)\cos\bar{\theta}_N + \sin\bar{\theta}_L(t)\sin\bar{\theta}_N\cos(\bar{\phi}_L(t) - \bar{\phi}_N)\,, \tag{2.33}$$

$$\hat{\mathbf{z}}\cdot\hat{\mathbf{n}} = \frac{1}{2}\cos\bar{\theta}_N - \frac{\sqrt{3}}{2}\sin\bar{\theta}_N\cos(\bar{\Phi}_D(t) - \bar{\phi}_N)\,, \tag{2.34}$$

$$\begin{aligned}\hat{\mathbf{n}}\cdot(\hat{\mathbf{L}}\times\hat{\mathbf{z}}) = {}& \frac{1}{2}\sin\bar{\theta}_L(t)\sin\bar{\theta}_N\sin(\bar{\phi}_L(t) - \bar{\phi}_N) \\ & - \frac{\sqrt{3}}{2}\cos\bar{\Phi}_D(t)(\cos\bar{\theta}_L(t)\sin\bar{\theta}_N\sin\bar{\phi}_N - \cos\bar{\theta}_N\sin\bar{\theta}_L(t)\sin\bar{\phi}_L(t)) \\ & - \frac{\sqrt{3}}{2}\sin\bar{\Phi}_D(t)(\cos\bar{\theta}_N\sin\bar{\theta}_L(t)\cos\bar{\phi}_L(t) - \cos\bar{\theta}_L(t)\sin\bar{\theta}_N\cos\bar{\phi}_N)\,.\end{aligned} \tag{2.35}$$

Here $\bar{\theta}_L(t)$ and $\bar{\phi}_L(t)$ are the angles describing the binary's orientation relative to the *barycenter* frame; they vary in time solely due to precession.

We now turn to the strain in detector II. Following Cutler, we construct the signal from detector II as

$$h_{\rm II}(t) = \frac{1}{\sqrt{3}}\left[h_{\rm I}(t) + 2h_{\rm II'}(t)\right]\,, \tag{2.36}$$

where $h_{\rm I}$ is the signal from detector I, (2.24), and $h_{\rm II'} = (\delta L_2(t) - \delta L_3(t))/L$ is the signal formed from the difference in the lengths of arms 2 and 3. This choice makes the noise in detector I uncorrelated with the noise in detector II; we will exploit this property in Sec. 2.3 to treat detectors I and II as independent detectors. From (2.36), we obtain

$$h_{\rm II}(t) = \frac{\sqrt{3}}{2}\left[\frac{1}{2}(h_{xy} + h_{yx})\right]\,. \tag{2.37}$$

The result is that detector II also behaves like a 90° interferometer (scaled by $\sqrt{3}/2$), but rotated by 45° with respect to detector I. Thus the antenna patterns for detector

II are

$$F_{\rm II}^{+}(\theta_N,\phi_N,\psi_N) = F_{\rm I}^{+}(\theta_N,\phi_N-\pi/4,\psi_N)\,, \tag{2.38}$$

$$F_{\rm II}^{\times}(\theta_N,\phi_N,\psi_N) = F_{\rm I}^{\times}(\theta_N,\phi_N-\pi/4,\psi_N)\,. \tag{2.39}$$

We now rewrite the waveform in terms of an amplitude and phase. Letting $i \in \{\mathrm{I},\mathrm{II}\}$ label detector number, the waveform as measured by detector i is

$$h_i(t) = 2\frac{\mathcal{M}^{5/3}(\pi f)^{2/3}}{D_L} A_{{\rm pol},i}(t)\cos[\Phi(t)+\varphi_{{\rm pol},i}(t)+\varphi_D(t)+\delta_p\Phi(t)]\,, \tag{2.40}$$

where

$$A_{{\rm pol},i}(t) = \frac{\sqrt{3}}{2}[(1+(\hat{\mathbf{L}}\cdot\hat{\mathbf{n}})^2)^2 F_i^{+}(t)^2 + 4(\hat{\mathbf{L}}\cdot\hat{\mathbf{n}})^2 F_i^{\times}(t)^2]^{1/2} \tag{2.41}$$

is the "polarization amplitude" (pictured in Fig. 2-1) and

$$\varphi_{{\rm pol},i}(t) = \tan^{-1}\left[\frac{2(\hat{\mathbf{L}}\cdot\hat{\mathbf{n}})F_i^{\times}(t)}{[1+(\hat{\mathbf{L}}\cdot\hat{\mathbf{n}})^2]F_i^{+}(t)}\right] \tag{2.42}$$

is the "polarization phase" [62]. We have also introduced the "Doppler phase" $\varphi_D(t)$, which arises from the detector's motion around the Sun and is given to lowest order by

$$\varphi_D(t) = 2\pi f(t) R_\oplus \sin\bar{\theta}_N \cos[\bar{\Phi}_D(t)-\bar{\phi}_N]\,, \tag{2.43}$$

where $R_\oplus = 1$ AU.

Much of our analysis is done in the frequency domain. We define the Fourier transform of the signal as

$$\tilde{h}(f) = \int_{-\infty}^{\infty} e^{2\pi i f t} h(t) dt\,. \tag{2.44}$$

To evaluate the Fourier transform, we make use of the stationary phase approximation [63, 194]. This approximation relies on the fact that the orbital time scale $T_{\rm orb}$ is much shorter than the precession time scale $T_{\rm prec}$, as well as the inspiral time scale $T_{\rm insp}$ and detector orbital time scale $T_D = 1$ yr. The result thus differs from the true Fourier

transform by terms of order $T_{\rm orb}/T_{\rm prec}$ and $T_{\rm orb}/T_{\rm insp}$ [10]. The Fourier transform is thus likely to be inaccurate near the end of the inspiral, when all of these time scales become comparable. Using (2.13) and (2.14), we have

$$\tilde{h}_i(f) = \sqrt{\frac{5}{96}}\frac{\pi^{-2/3}\mathcal{M}^{5/6}}{D_L}A_{{\rm pol},i}[t(f)]f^{-7/6}e^{i(\Psi(f)-\varphi_{{\rm pol},i}[t(f)]-\varphi_D[t(f)]-\delta_p\Phi[t(f)])}\,, \qquad (2.45)$$

where the phase $\Psi(f)$ is given by

$$\begin{aligned}\Psi(f) = 2\pi f t_c - \Phi_c - \frac{\pi}{4} + \frac{3}{128}(\pi\mathcal{M}f)^{-5/3}\Bigg[1 + \frac{20}{9}\left(\frac{743}{336}+\frac{11}{4}\eta\right)(\pi Mf)^{2/3} \\ -4(4\pi-\beta)(\pi Mf) + 10\left(\frac{3058673}{1016064}+\frac{5429}{1008}\eta+\frac{617}{144}\eta^2-\sigma\right)(\pi Mf)^{4/3}\Bigg]\,.\end{aligned} \qquad (2.46)$$

In the work by Cutler [62], the separation of time scales that we used above leads to an interpretation of the polarization amplitude, polarization phase, and Doppler phase as modulations, in amplitude and phase, of an underlying carrier signal. These modulations make it possible to measure the sky position of the source, which also helps to measure the luminosity distance D_L [127]. With the addition of precession, the polarization amplitude and polarization phase include additional modulations which further improve the measurement of these parameters. In conjunction with the other effects of precession (on β, σ, and $\delta_p\Phi[t(f)]$), these effects also help us to better measure the masses and spins of the system.

2.3 Measurement and parameter estimation with LISA

2.3.1 Theory

In the previous section, we constructed the expected form for the GW strain that LISA is being designed to measure. The signal $s_i(t)$ as measured by detector i will of

course also include noise $n_i(t)$:

$$s_i(t) = h_i(t) + n_i(t)\,. \tag{2.47}$$

The LISA noise spectrum is discussed in section 2.3.2; in this section, we discuss the theory of parameter estimation with a noisy signal. First, consider only one detector. We assume that the noise is zero mean, wide-sense stationary, and Gaussian. Wide-sense stationary means that the autocovariance function

$$K_n(t,t') = \langle n(t)n(t')\rangle - \langle n(t)\rangle\langle n(t')\rangle \tag{2.48}$$

depends only on the time difference $\tau = t - t'$. (Throughout this section, quantities within angle brackets are ensemble averaged with respect to the noise distribution.) A process is Gaussian if every sample of the process can be described as a Gaussian random variable and all possible sets of samples of the process are jointly Gaussian. However, the noise is colored, not white. A white noise process is defined to be a process which is uncorrelated with itself at different times; that is, its autocovariance is a delta function. Because the noise is colored, it has an interesting (nonflat) power spectral density (PSD), which is defined as the Fourier transform of the autocovariance function:

$$S_n(f) = 2\int_{-\infty}^{\infty} d\tau e^{2\pi i f\tau} K_n(\tau)\,. \tag{2.49}$$

The factor of 2 follows [63]; we actually use the *one-sided* PSD. Since the noise is Gaussian, it is described entirely by its second moments. Therefore, we will only need the PSD, and not the full probability density function, to analyze the effect of the noise on the signal.

Incidentally, it can be shown that wide-sense stationarity implies that the Fourier transform of $n(t)$ is a nonstationary white noise process in frequency:

$$\langle \tilde{n}(f)\tilde{n}^*(f')\rangle = \frac{1}{2}\delta(f-f')S_n(f)\,. \tag{2.50}$$

The Fourier components are thus independent Gaussian random variables.

Now briefly consider both detectors. We explicitly constructed the second detector (2.36) (with $h(t) \to s(t)$) so that the noise in it is uncorrelated with, and thus independent of, noise in the first detector. Thus we have

$$\langle \tilde{n}_i(f) \tilde{n}_j^*(f') \rangle = \frac{1}{2} \delta_{ij} \delta(f - f') S_n(f) \,. \tag{2.51}$$

The uncorrelated nature of these two noises will allow us to easily generalize discussion from one detector to the full two effective detector system.

Let us write our GW as $h(\boldsymbol{\theta})$, where the components of the vector $\boldsymbol{\theta}$ represent the various parameters on which the waveform depends. We now assume that a GW signal with particular parameters $\tilde{\boldsymbol{\theta}}$ is present in the data (i.e., "detection" has already occurred), and want to obtain estimates $\hat{\boldsymbol{\theta}}$ of those source parameters. Finn [90] shows that the probability for the noise to have some realization $n_0(t)$ is given by

$$p(n = n_0) \propto e^{-(n_0|n_0)/2} \,, \tag{2.52}$$

where the inner product used here is given by

$$(a|b) = 4 \,\mathrm{Re} \int_0^\infty df \frac{\tilde{a}^*(f) \tilde{b}(f)}{S_n(f)} \tag{2.53}$$

$$= 2 \int_0^\infty df \frac{\tilde{a}^*(f) \tilde{b}(f) + \tilde{a}(f) \tilde{b}^*(f)}{S_n(f)} \,. \tag{2.54}$$

This product is a natural one for the vector space of (frequency-domain) signals $a(f)$. (Note that this definition of the inner product differs from [90] by a factor of 2.)

Given a particular measured signal $s(t)$, the probability that the GW parameters are given by $\tilde{\boldsymbol{\theta}}$ is the same as the probability that the noise takes the realization $s - h(\tilde{\boldsymbol{\theta}})$:

$$p(\tilde{\boldsymbol{\theta}}|s) \propto e^{-(h(\tilde{\boldsymbol{\theta}})-s|h(\tilde{\boldsymbol{\theta}})-s)/2} \,, \tag{2.55}$$

where the constant of proportionality may include prior probability densities for the parameters $\tilde{\boldsymbol{\theta}}$. For simplicity, we take these to be uniform.

We can estimate the parameters $\tilde{\boldsymbol{\theta}}$ by the maximum likelihood (ML) method. This method involves finding the parameters $\hat{\boldsymbol{\theta}}$ that maximize (2.55)[9], or alternatively, minimize $(h(\tilde{\boldsymbol{\theta}}) - s|h(\tilde{\boldsymbol{\theta}}) - s)$, which can be considered a distance in signal space. A bank of template waveforms is correlated with the received signal and, assuming that any template produces a statistically significant correlation, the one with the highest correlation is the one with the ML parameters. The SNR for this signal is then given by [63]

$$\rho \approx (h(\hat{\boldsymbol{\theta}})|h(\hat{\boldsymbol{\theta}}))^{1/2} \approx (h(\tilde{\boldsymbol{\theta}})|h(\tilde{\boldsymbol{\theta}}))^{1/2} \,. \tag{2.56}$$

To quantify the errors in the ML estimate, we expand (2.55) around the most likely values $\hat{\boldsymbol{\theta}}$. We can then write the probability density as [63, 194]:

$$p(\tilde{\boldsymbol{\theta}}|s) \propto e^{-\Gamma_{ab}\delta\theta^a\delta\theta^b/2} \,, \tag{2.57}$$

where $\delta\theta^a = \tilde{\theta}^a - \hat{\theta}^a$ and

$$\Gamma_{ab} = \left(\frac{\partial h}{\partial \theta^a} \middle| \frac{\partial h}{\partial \theta^b} \right) \,, \tag{2.58}$$

evaluated at $\boldsymbol{\theta} = \hat{\boldsymbol{\theta}}$, is known as the *Fisher information matrix*. For small deviations from the ML estimate, the distribution is Gaussian. This expression holds for large values of the SNR (2.56). It is worth emphasizing at this point that, in our evaluation of 2.58, most derivatives are taken numerically using finite differencing — the complicated nature of the signal (due to the inclusion of spin precession) makes it essentially impossible to evaluate all but a few of our derivatives analytically. This is another reason that the code we have developed for this analysis is substantially slower than those developed for analyses which do not include spin-precession physics.

Now we return again to the two detector case. Using (2.51), we can write a total Fisher matrix as the sum of the individual Fisher matrices for each detector:

$$\Gamma^{\rm tot}_{ab} = \Gamma^{\rm I}_{ab} + \Gamma^{\rm II}_{ab} \,. \tag{2.59}$$

[9]Throughout this chapter, we refer to (2.55) as the "likelihood," when technically it is the *posterior probability density*. The likelihood is properly defined as $p(s|\tilde{\boldsymbol{\theta}})$. For the case of uniform priors, the two quantities are equivalent up to a normalization, and this terminology, while sloppy, is acceptable.

The Fisher matrix is then inverted to produce the covariance matrix $\Sigma^{ab} = (\Gamma^{-1}_{\rm tot})^{ab}$. The diagonal terms of the covariance matrix represent measurement errors:

$$\Delta\theta^a \equiv \sqrt{\langle(\delta\theta^a)^2\rangle} = \sqrt{\Sigma^{aa}}\,. \tag{2.60}$$

The off-diagonal terms can be expressed as correlation coefficients, ranging from -1 to 1:

$$c^{ab} \equiv \frac{\langle\delta\theta^a\delta\theta^b\rangle}{\Delta\theta^a\Delta\theta^b} = \frac{\Sigma^{ab}}{\sqrt{\Sigma^{aa}\Sigma^{bb}}}\,. \tag{2.61}$$

2.3.2 LISA detector and astrophysical noise

We turn now to a discussion of the noise we expect in LISA measurements. Our model for the instrumental noise spectrum, $S_h^{\rm inst}(f)$, is based on that described in [159]. (From now on, we use the notation S_h for strain noise instead of S_n for general noise.) In particular, we use the online sensitivity curve generator provided by Shane Larson[10], which implements the recipe of [159]. The output of Larson's webtool gives a sky averaged *amplitude* sensitivity curve, $h_{\rm Larson}$. To convert to the noise we need for our analysis, we square this amplitude and insert two numerical factors:

$$S_h^{\rm inst}(f) = \frac{1}{5}\times\left(\frac{\sqrt{3}}{2}h_{\rm Larson}\right)^2 = \frac{3}{20}h^2_{\rm Larson}\,. \tag{2.62}$$

The factor of $1/5$ accounts for the averaging of the antenna pattern functions over all sky positions and source orientations. This factor is only correct for measuring radiation with wavelength $\lambda \gg L$ (where L is the LISA arm length). As a consequence, our instrumental noise will be inaccurate at high frequencies. This will have little impact on our analysis since, as already argued, the signal from merging binary black holes accumulates at low frequencies.

The factor $\sqrt{3}/2$ arises due to the 60° opening angle of the interferometer arms; we have already accounted for this factor in our discussion of the interferometer's interaction with a GW (cf. (2.25) and (2.37)). The numerical factor 3/20 has been

[10]This can be found at `http://www.srl.caltech.edu/~shane/sensitivity/`.

the source of some confusion; Berti, Buonanno and Will very nicely straightened this out. See Sec. IIC of [29] for further discussion of these factors.

Besides purely instrumental noise, LISA data will contain "noise" from a background of confused binary sources[11], mostly white dwarf binaries. An isotropic background of indistinguishable sources can be represented as noise with spectral density [21]

$$S_h^{\rm conf}(f) = \frac{3}{5\pi} f^{-3} \rho_c \Omega_{\rm GW}(f)\,, \tag{2.63}$$

where $\rho_c = 3H_0^2/8\pi$ is the critical energy density to close the universe and $\Omega_{\rm GW} = (f/\rho_c) d\rho_{\rm GW}/df$ is the energy density in GWs relative to ρ_c per logarithmic frequency interval. Using this form and the results of Farmer and Phinney [84], we model the confusion noise due to extragalactic binary sources by

$$S_h^{\rm exgal}(f) = 4.2 \times 10^{-47} \left(\frac{f}{1\,{\rm Hz}}\right)^{-7/3} {\rm Hz}^{-1}\,. \tag{2.64}$$

From Nelemans et al. [183], we take the galactic white dwarf confusion noise to be

$$S_h^{\rm gal}(f) = 2.1 \times 10^{-45} \left(\frac{f}{1\,{\rm Hz}}\right)^{-7/3} {\rm Hz}^{-1}\,. \tag{2.65}$$

The combined instrumental and galactic confusion noise is given by [21]

$$S_h^{\rm inst+gal}(f) = \min[S_h^{\rm inst}(f)/\exp(-\kappa T_{\rm mission}^{-1} dN/df), S_h^{\rm inst}(f) + S_h^{\rm gal}(f)]\,. \tag{2.66}$$

The choice taken in (2.66) reflects the fact that, at sufficiently high frequency, the number of binaries per bin should be small enough that they are no longer truly confused and can be subtracted from the data stream (at least partially). The factor $\exp(-\kappa T_{\rm mission}^{-1} dN/df)$ is the fraction of "uncorrupted" frequency bins. We choose $\kappa = 4.5$ [55], $T_{\rm mission}$ is the mission duration (which we take to be three years), and

$$\frac{dN}{df} = 2 \times 10^{-3} \left(\frac{1\,{\rm Hz}}{f}\right)^{11/3} {\rm Hz}^{-1} \tag{2.67}$$

[11] While surely noise when studying cosmological black holes, this background is *signal* to those interested in stellar populations.

is the number density of galactic binaries per unit frequency [127].

Finally, the total noise is given by

$$S_h(f) = S_h^{\rm inst+gal}(f) + S_h^{\rm exgal}(f)\,. \tag{2.68}$$

2.4 Results

2.4.1 Procedural issues

Parameter space

Seventeen parameters describe the most general binary black hole inspiral waveform [250]. Two of these are the orbital eccentricity and the orientation of the orbital ellipse; since we only consider circular orbits, we can ignore these two. The other 15 parameters are all necessary to describe the full post-Newtonian waveform with precession effects that we described in section 2.2.

We divide this set into intrinsic and extrinsic parameters. In our system, intrinsic parameters are those which label properties intrinsic to the binary itself; extrinsic parameters label properties which depend upon the position and placement of the binary relative to the observer. One can regard intrinsic parameters as describing the physics or astrophysics of the binary system, and extrinsic parameters as describing the binary's astronomical properties.

The intrinsic parameters we use are $\ln m_1$; $\ln m_2$; χ_1 and χ_2, the dimensionless spin parameters; $\bar{\mu}_{S_1}(0) \equiv \cos[\bar{\theta}_{S_1}(0)]$, $\bar{\phi}_{S_1}(0)$, $\bar{\mu}_{S_2}(0) \equiv \cos[\bar{\theta}_{S_2}(0)]$, and $\bar{\phi}_{S_2}(0)$, the initial directions of the spins; t_c, the time at coalescence; and Φ_c, the phase at coalescence. (Note that t_c and Φ_c could very well be considered extrinsic, since they just label the system's state at some particular time. At any rate, neither t_c nor Φ_c is of much physical interest, so their categorization is not too important.) Our extrinsic parameters are $\bar{\mu}_L(0) \equiv \cos[\bar{\theta}_L(0)]$ and $\bar{\phi}_L(0)$, the initial direction of the orbital angular momentum; $\bar{\mu}_N = \cos\bar{\theta}_N$ and $\bar{\phi}_N$, the sky position in barycenter coordinates; and $\ln D_L$, the luminosity distance to the binary. All of these parameters must be

fit in a measurement and thus must be included in our Fisher matrix analysis. We are not necessarily interested in all of them, however. In particular, we will focus on the masses, the dimensionless spin parameters, the sky position, and the luminosity distance.

It is worth noting that this choice of parameters is not the same as that used in analyses which neglect precession. In that case, the direction of the angular momentum $\hat{\mathbf{L}}$ is constant and fully described using two numbers (e.g., $\bar{\mu}_L$ and $\bar{\phi}_L$). Including precession, $\hat{\mathbf{L}}$ is no longer constant, but evolves according to (2.17). The solution to this differential equation requires two initial conditions, for instance, $\bar{\mu}_L(0)$ and $\bar{\phi}_L(0)$, which can be used as parameters of the system. Since these initial conditions are taken at the (somewhat arbitrary) starting point of our calculations, they do not hold much physical interest (though they must be fit for and thus included in our Fisher matrix).

Previous analyses, including the precursor to this work [127], have used β (2.11) and σ (2.12) as parameters — these are constants when precession is neglected. They are also the only combinations of the spin magnitudes and spin angles that enter into the expression for the waveform. Boiling the six numbers which characterize $\mathbf{S}_1$ and $\mathbf{S}_2$ down to two greatly simplifies the parameter space, but also restricts us from being able to measure, for example, the black holes' spin magnitudes. When precession is included, β and σ are no longer constants. In addition, they no longer fully characterize the signal, since the precession equations (2.15), (2.16), and (2.17) depend on all of the components of the spins. We thus need six spin-related parameters to fully describe the signal: the magnitudes of the spins and their orientations at some initial time. The orientations are again uninteresting, but the fact that we can measure the magnitudes of the spins and quantify their errors is quite interesting and new to this analysis.

Finally, we break from tradition and use $\ln m_1$ and $\ln m_2$ to parameterize our masses rather than $\ln \mathcal{M}$ and $\ln \mu$. The chirp mass and reduced mass have been used in most previous work because of their appearance in the waveform phase $\Psi(f)$. However, the precession equations, as well as the spin parameters β and σ, depend

on the individual masses of the black holes. It is a simple matter in principle to just solve for $m_{1,2}(\mathcal{M},\mu)$ and substitute into the precession equations. Unfortunately, the Jacobian of the transformation between $(\mathcal{M},\mu)$ and (m_1,m_2) is singular when $m_1 = m_2$, leading to problems in evaluating the Fisher matrix.

These problems can be illustrated analytically. Consider how derivatives of some function $f(m_1,m_2)$ with respect to $\mathcal{M}$ behave:

$$\frac{\partial f}{\partial \mathcal{M}} = \sum_{i=1}^{2} \frac{\partial f(m_1,m_2)}{\partial m_i}\frac{\partial m_i(\mathcal{M},\mu)}{\partial \mathcal{M}}\,. \tag{2.69}$$

When $m_1 = m_2$, the second of these derivatives diverges — a behavior that we have seen numerically. The Fisher information is infinite, and the Gaussian approximation breaks down; the same problem occurs for μ. Thus, we argue that, when precession is included, $\mathcal{M}$ and μ are no longer a good choice of parameters to describe the system. Since we are still interested in the errors in $\ln\mathcal{M}$ and $\ln\mu$ (which are determined to higher accuracy than the individual masses), we convert using the propagation of errors formulas

$$\begin{aligned}\left(\frac{\Delta\mathcal{M}}{\mathcal{M}}\right)^2 = &\left(\frac{m_1}{\mathcal{M}}\right)^2\left(\frac{\partial\mathcal{M}}{\partial m_1}\right)^2\left(\frac{\Delta m_1}{m_1}\right)^2 + \left(\frac{m_2}{\mathcal{M}}\right)^2\left(\frac{\partial\mathcal{M}}{\partial m_2}\right)^2\left(\frac{\Delta m_2}{m_2}\right)^2 \\ &+ 2\left(\frac{m_1 m_2}{\mathcal{M}^2}\right)\left(\frac{\partial\mathcal{M}}{\partial m_1}\right)\left(\frac{\partial\mathcal{M}}{\partial m_2}\right)\Sigma^{\ln m_1,\ln m_2}\,,\end{aligned} \tag{2.70}$$

$$\begin{aligned}\left(\frac{\Delta\mu}{\mu}\right)^2 = &\left(\frac{m_1}{\mu}\right)^2\left(\frac{\partial\mu}{\partial m_1}\right)^2\left(\frac{\Delta m_1}{m_1}\right)^2 + \left(\frac{m_2}{\mu}\right)^2\left(\frac{\partial\mu}{\partial m_2}\right)^2\left(\frac{\Delta m_2}{m_2}\right)^2 \\ &+ 2\left(\frac{m_1 m_2}{\mu^2}\right)\left(\frac{\partial\mu}{\partial m_1}\right)\left(\frac{\partial\mu}{\partial m_2}\right)\Sigma^{\ln m_1,\ln m_2}\,.\end{aligned} \tag{2.71}$$

For unequal masses, we find that computing errors in m_1 and m_2 and then converting gives the same result as simply computing errors in $\mathcal{M}$ and μ directly. We do not find good agreement in the equal-mass case; for the reasons discussed above, however, we do not trust the $(\mathcal{M},\mu)$ parameterization in this case. At any rate, the case $m_1 = m_2$ is quite implausible in nature, so this is almost certainly a moot point as far as real

measurements are concerned.[12] We note that Vecchio [250], for simplicity, considers the equal-mass case exclusively but does not report any anomalous behavior such as we have seen. We are puzzled about this discrepancy. On the other hand, similar behavior is seen by Trias and Sintes [245], who ignore precession but do include higher harmonics in the waveform model (i.e., they go beyond the restricted post-Newtonian approximation).

Calculations

The code we use to calculate parameter measurement errors is based on that used in [127]. It is written in C++ using several routines taken, sometimes with slight modification, from [200]. As in [127], we perform Monte Carlo simulations in which we specify rest-frame masses and redshift and then randomly choose the sky position, initial angular momentum direction, and initial spin directions for each binary. In some cases, we specify spin magnitudes, but in most cases, we choose them randomly as well. We also uniformly distribute the time parameter t_c of each binary over the assumed duration of the LISA mission (which we take to be 3 yr).

The primary function of the code is the calculation of the full gravitational waveform, including precession effects. In order to effectively use the formulas of section 2.3, we take the wave frequency f as the independent variable. The elapsed time is related to the frequency using (2.13). The calculation is started when the waveform enters LISA's band (taken to be $f_{\rm min} = 3 \times 10^{-5}$ Hz throughout this thesis) or when the LISA mission begins, whichever is later. By treating the time of coalescence as a Monte Carlo variable, some signals will be partially cut off because they are already in band when LISA begins observations.

The end of inspiral/beginning of merger is a somewhat ad hoc and fuzzy boundary. Indeed, recent numerical computations have shown that the GWs produced by a binary that coalesces into a single body do not show any particular special feature as the black holes come together, instead smoothly chirping through this transition [44, 186]. Since we are not including numerical merger waves in our analysis, we require

[12] Even a slight mass difference (a few percent) is sufficient for the two approaches to match.

some point to terminate our post-Newtonian expansion. Most studies show that the inspiral comes to an end when the separation of the bodies in harmonic coordinates is roughly $r \sim 6M$; at this point, the system's GW frequency is approximately given by

$$f_{\rm merge} \simeq \frac{2}{2\pi}\Omega_{\rm Kepler}(r = 6M) \simeq \frac{0.02}{M}\,, \tag{2.72}$$

where $\Omega_{\rm Kepler} = (M/r^3)^{1/2}$ is the Keplerian orbital angular frequency. (The factor $1/2\pi$ converts from angular frequency to frequency; the additional factor of 2 accounts for the quadrupolar nature of gravitational waves.) We use equation (2.72) throughout our analysis to terminate the inspiral.

Once the frequency range has been determined, the true work begins. We integrate the precession equations (2.15), (2.16), and (2.17) using a Runge-Kutta routine to find the values of $\hat{\mathbf{L}}$, $\hat{\mathbf{S}}_1$, and $\hat{\mathbf{S}}_2$ over the duration of the signal. The routine is a fifth-order adaptive-step algorithm in the frequency domain. At each frequency, the code takes the results for the three orbital angular momentum components and six spin components and uses them to calculate $\bar{\mu}_L$, $\bar{\phi}_L$, β, and σ. It also computes the integrated correction to the phase using the derivative (2.23).

As already discussed, our derivatives are taken numerically rather than analytically. We therefore must do the integration described above a total of 21 times: once for the given values of the parameters, and twice more for small shifts in each parameter which requires a numerical derivative. This repetition slows the code quite drastically compared to its earlier incarnation — an unfortunate but unavoidable cost.

Once all of the necessary integrations are complete, the SNR (2.56) and the Fisher matrix (2.58) can be calculated for each of the two effective detectors of LISA using the noise $S_h(f)$ (see Sec. 2.3.2). Some previous work [62, 29] investigated parameter estimation using the signal from only one synthesized detector; we will always assume that both are operational. It would be interesting to see how measurement degradation due to only having a single operating detector can be ameliorated by including precession effects.

At this stage, the necessary integrals are performed using Curtis-Clenshaw quadrature, which depends on the decomposition of the integrand into Chebyshev polynomials [200]. This method keeps the code reasonably fast even with the addition of the Runge-Kutta routine. At each step of the integration, the integrator uses the values that were calculated using that Runge-Kutta routine to evaluate the waveform and/or its appropriate derivatives. The derivatives are calculated using

$$\frac{df}{d\theta} \approx \frac{f(\theta + \frac{\Delta\theta}{2}) - f(\theta - \frac{\Delta\theta}{2})}{\Delta\theta} \,. \tag{2.73}$$

For all parameters, we use $\Delta\theta = 10^{-5}\,\theta$. We invert the Fisher matrix using LU decomposition to produce the covariance matrix [200]. In "poor" cases (e.g., high-mass binaries at large redshift), the Fisher matrix can be nearly singular, with a large condition number.[13] In such a case, the covariance matrix produced by the code may not be the true inverse of the Fisher matrix (and may not even be positive definite). This problem is largely ameliorated by representing our numerical data in `long double` format — this improves (relative to type `double`) matrix inverses in many "bad" cases but leaves all other cases essentially unchanged.

It is worth noting that the bad cases are typically ones in which the binary executes very few orbits over the course of the measurement. We are confident in our results for all cases in which the number of measured orbits, $N_{\rm orb} = \Phi_{\rm orb}/2\pi$, is greater than $\sim 10 - 20$. When the number of orbits is small (and the condition number is concomitantly high), the errors are so large that they are basically meaningless. In such a case, measurement would not determine the system's characteristics in any meaningful sense.

[13]The "condition number" is the ratio of the largest eigenvalue of a matrix to the smallest. A rule of thumb is that matrix inversion breaks down when the logarithm of the condition number of a matrix exceeds the number of digits of accuracy in the matrix elements (see, e.g., discussion in [200]).

2.4.2 Black hole masses and spins

Representative examples of our results are shown in Figures 2-2 and 2-3. These histograms show the spread of errors in $\mathcal{M}$ and μ for a sample of 10^4 binaries at $z = 1$ with rest frame masses $m_1 = 10^6\,M_\odot$ and $m_2 = 3\times10^5\,M_\odot$. Each figure compares the results of the new code to those of the original code of [127], which neglects precession. (That code has been updated to reflect up-to-date models for LISA noise; some minor coding errors and one major one [155] have also been corrected.) Clearly, including spin precession leads to a significant improvement in the measurement of these mass parameters. The reduced mass μ, in particular, is improved. This is because the time variation of β and σ breaks a near degeneracy between those terms and μ in the post-Newtonian phase (2.46). The masses also control the precession rate, as seen in (2.15), (2.16), and (2.17). (Recall that, in those equations, $S_i = \chi_i m_i^2$.) This means that they now influence the polarization amplitude and polarization phase; they do not influence those quantities when precession is neglected. These precession-induced influences on the waveform make it possible to determine the masses even more accurately than before.

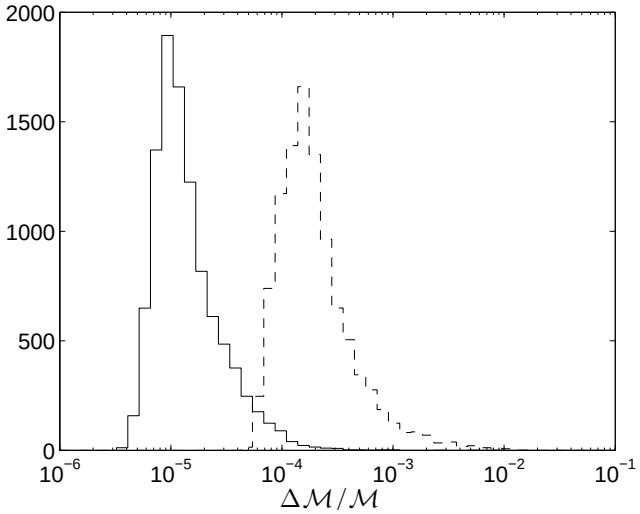

Figure 2-2: Distribution of errors in chirp mass $\mathcal{M}$ for 10^4 binaries with $m_1 = 10^6 M_\odot$ and $m_2 = 3\times10^5 M_\odot$ at $z = 1$. The dashed line is the precession-free calculation; the solid line includes precession. Precession reduces the measurement error by about an order of magnitude.

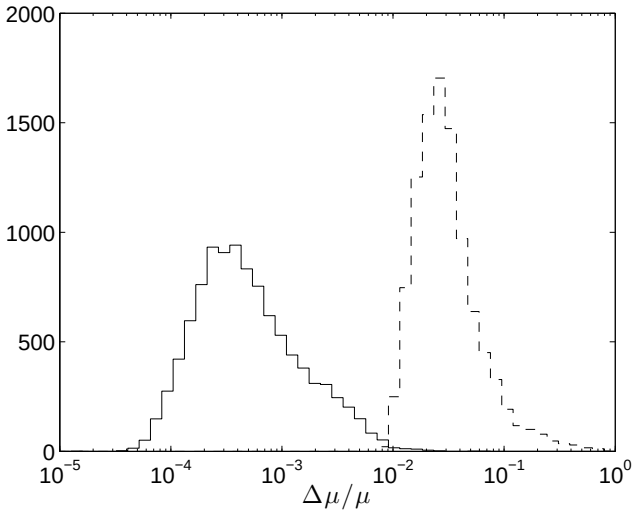

Figure 2-3: Distribution of errors in reduced mass μ for 10^4 binaries with $m_1 = 10^6 M_\odot$ and $m_2 = 3 \times 10^5 M_\odot$ at $z = 1$. The dashed line is the precession-free calculation; the solid line includes precession. Precession has an enormous effect on the reduced mass, which was previously highly correlated with the parameters β and σ.

As discussed earlier, we have found the masses m_1 and m_2 to be more useful parameters than $\mathcal{M}$ and μ when precession is included. Figure 2-4 shows the error in measurements of the individual masses for our example system. While these masses are measured quite accurately, they are not measured as accurately as $\mathcal{M}$ and μ. This reflects the fact that, even though the individual masses play a role in the precession, the other parts of the waveform depend explicitly on the combinations $\mathcal{M}$ and μ. Notice also that the smaller mass is typically determined a bit better than the larger one, though the difference is not large.

Precession makes it possible to determine the spins of the binary's members. Figure 2-5 shows the error in measurements of the two dimensionless spin parameters χ_1 and χ_2. We see that χ is generally determined very well: Taking a typical spin parameter to be about 0.5 (recall we randomly choose χ between 0 and 1), the bulk of this distribution corresponds to errors of a bit less than a percent. For this entirely random distribution of χ, the dimensionless spin parameter of the larger hole tends to be better determined than that of the smaller hole. This appears to be a simple consequence of the fact that black hole spin scales as mass squared ($S_i = \chi_i m_i^2$), and

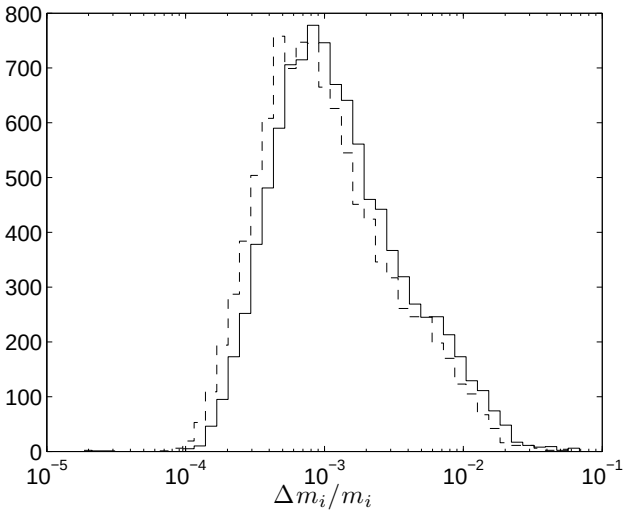

Figure 2-4: Distribution of errors in individual hole masses for 10^4 binaries at $z = 1$. The solid line is $m_1 = 10^6 M_\odot$, while the dashed line is $m_2 = 3 \times 10^5 M_\odot$. The individual masses are not determined as well as $\mathcal{M}$ and μ, but they are better behaved parameters when precession is introduced.

larger spin has more of an impact on the waveform.

Next, we examine how well spin is measured as a function of spin magnitude. Figure 2-6 shows the error in χ_1 for the same system as in Fig. 2-5, except that we set $\chi_1 = \chi_2 = 0.9$ (solid line) and $\chi_1 = \chi_2 = 0.1$ (dashed line), rather than randomly distributing their values. This allows us to more accurately assess how well spin is determined as a function of its value, as well as to more accurately determine the percent error we expect in these measurements. For $\chi_1 = \chi_2 = 0.1$, the error is almost 10%, while for $\chi_1 = \chi_2 = 0.9$, the error is closer to 0.1%. This is a considerable difference and is easily ascribed to the fact that rapid spin has a much stronger impact on the waveform.

Table 2.1 shows the median errors in intrinsic parameters for different masses at $z = 1$. We continue to include the errors in $\mathcal{M}$ and μ for comparison with the precession-free case, but only in binaries of unequal mass where the Gaussian approximation is well defined. Examining the table, we see some interesting features. The errors, in general, are worse for higher-mass binaries, which spend less time in the LISA band. At $m_1 = m_2 = 10^7\, M_\odot$, the mass errors jump to nearly 10%, compared

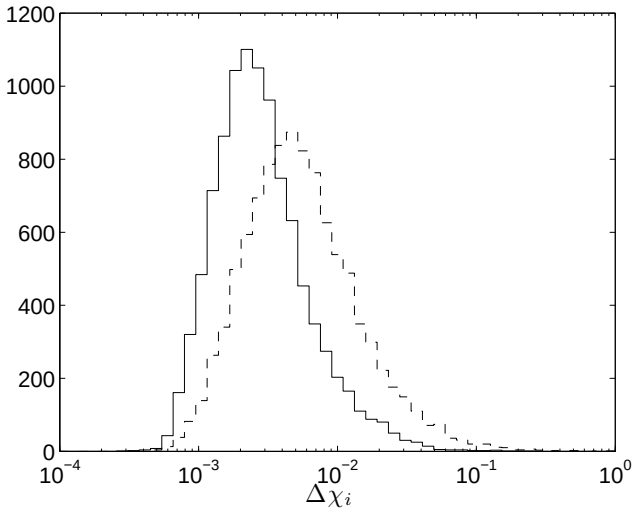

Figure 2-5: Distribution of errors in dimensionless spin parameters χ_1 (solid line) and χ_2 (dashed line) for 10^4 binaries with $m_1 = 10^6 M_\odot$ and $m_2 = 3 \times 10^5 M_\odot$ at $z = 1$. In each binary, the spin values are randomly selected between 0 and 1. The higher mass then has, on average, higher total spin and more effect on the precession.

to tenths of a percent at the next lower mass combination. In addition, the spin determination becomes very unreliable. Mass ratio also has an important effect on the results. Taking into account the general trend caused by total mass, we see that nonunity mass ratios generally produce better results. This is good news for eventual measurements of astrophysical systems, since merger tree calculations show that binaries are most likely to have mass ratios of about 10 [219]. To understand the mass ratio dependence, we again turn to the precession equations (2.15), (2.16), and (2.17). For unequal masses, the geodetic spin-orbit and spin-spin terms will cause the two spins to precess at different rates, creating richer features in the signal than for equal masses. This illustrates the importance of effects beyond the "simple precession" of [10, 250]. We also see that the trends of Figs. 2-4 and 2-5 hold for each unequal-mass binary in the table. That is, the mass of the smaller hole is determined better than the mass of the larger hole, but the spin of the larger hole is determined better than the spin of the smaller hole.

Tables 2.2 and 2.3 show the same results for $z = 3$ and $z = 5$, respectively. The trends we see at $z = 1$ largely continue at these redshifts. In general, the

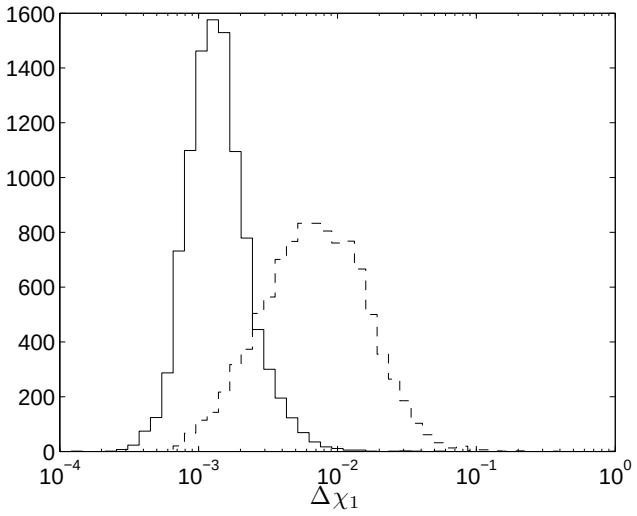

Figure 2-6: Distribution of errors in dimensionless spin parameter χ_1 for 10^4 binaries with $m_1 = 10^6 M_\odot$ and $m_2 = 3 \times 10^5 M_\odot$ at $z = 1$. Here, spin magnitudes have been set to a specified value — low spin, $\chi_1 = \chi_2 = 0.1$ (dashed line), and high spin, $\chi_1 = \chi_2 = 0.9$ (solid line). Since greater spin more strongly impacts the waveform, the high spin case is measured more accurately.

errors get worse at higher redshift as the signal amplitude degrades and more of the signal is redshifted out of band. It is worth noting that the change is generally greater from $z = 1$ to $z = 3$ than from $z = 3$ to $z = 5$. This effect was also seen by Berti, Buonanno, and Will [29] and can be explained by considering the redshift dependence of the wave amplitude. Neglecting all the angular factors and remembering to redshift quantities with the dimensions of time, we find that the amplitude scales like $(1 + z)/D_L(z) = 1/D_M(z)$, where $D_M(z)$ is the proper motion distance. This distance measure varies more strongly with z at low redshift than at high redshift. (See [120] for a plot of $D_M(z)$.) Consequently, when moving from $z = 1$ to $z = 3$, the amplitude, and thus the SNR, decreases more than when moving from $z = 3$ to $z = 5$. For lower-mass binaries, this amplitude decrease plays a bigger role in the loss of SNR than does redshifting the spectrum to lower frequency; most of the SNR is accumulated late in the inspiral, where the orbits are in a relatively flat region of the sensitivity curve.

By contrast, for the highest-mass binaries, redshifting of the spectrum can have

m_1 ($M_\odot$)	m_2 ($M_\odot$)	$\Delta m_1/m_1$	$\Delta m_2/m_2$	$\Delta\chi_1$	$\Delta\chi_2$	$\Delta\mathcal{M}/\mathcal{M}$ (no precession)	$\Delta\mathcal{M}/\mathcal{M}$ (precession)	$\Delta\mu/\mu$ (no precession)	$\Delta\mu/\mu$ (precession)
10^5	10^5	0.000783	0.000782	0.00415	0.00414	—	—	—	—
3×10^5	10^5	0.000667	0.000541	0.00157	0.00306	5.92×10^{-5}	5.51×10^{-6}	0.0114	0.000239
3×10^5	3×10^5	0.00109	0.00109	0.00539	0.00536	—	—	—	—
10^6	10^5	0.000629	0.000440	0.00102	0.00440	0.000156	1.18×10^{-5}	0.0180	0.000343
10^6	3×10^5	0.00111	0.000882	0.00256	0.00499	0.000170	1.19×10^{-5}	0.0274	0.000423
10^6	10^6	0.00195	0.00195	0.00902	0.00897	—	—	—	—
3×10^6	3×10^5	0.000988	0.000691	0.00137	0.00563	0.000583	2.53×10^{-5}	0.0550	0.000539
3×10^6	10^6	0.00238	0.00192	0.00380	0.00674	0.00117	4.19×10^{-5}	0.135	0.000849
3×10^6	3×10^6	0.00584	0.00582	0.0271	0.0275	—	—	—	—
10^7	10^6	0.00239	0.00177	0.00233	0.0122	0.00770	0.000174	0.469	0.00140
10^7	3×10^6	0.00814	0.00671	0.00829	0.0159	0.00851	0.000436	0.607	0.00332
10^7	10^7	0.0804	0.0802	0.492	0.493	—	—	—	—

Table 2.1: Median errors in intrinsic quantities for 10^4 binaries of various masses at $z = 1$, including comparisons with the "no precession" case where possible. We have omitted the errors in chirp mass and reduced mass for equal-mass binaries because that parameterization of the waveform fails the Gaussian approximation at those points.

m_1 ($M_\odot$)	m_2 ($M_\odot$)	$\Delta m_1/m_1$	$\Delta m_2/m_2$	$\Delta\chi_1$	$\Delta\chi_2$	$\Delta\mathcal{M}/\mathcal{M}$ (no precession)	$\Delta\mathcal{M}/\mathcal{M}$ (precession)	$\Delta\mu/\mu$ (no precession)	$\Delta\mu/\mu$ (precession)
10^5	10^5	0.00362	0.00362	0.0187	0.0185	—	—	—	—
3×10^5	10^5	0.00363	0.00294	0.00879	0.0171	0.000406	3.31×10^{-5}	0.0715	0.00130
3×10^5	3×10^5	0.00569	0.00569	0.0271	0.0269	—	—	—	—
10^6	10^5	0.00330	0.00231	0.00498	0.0208	0.00120	7.09×10^{-5}	0.128	0.00180
10^6	3×10^5	0.00648	0.00517	0.0120	0.0229	0.00174	9.17×10^{-5}	0.228	0.00248
10^6	10^6	0.0138	0.0139	0.0627	0.0630	—	—	—	—
3×10^6	3×10^5	0.00569	0.00402	0.00664	0.0287	0.00633	0.000241	0.456	0.00314
3×10^6	10^6	0.0181	0.0148	0.0223	0.0386	0.00708	0.000554	0.596	0.00658
3×10^6	3×10^6	0.0744	0.0737	0.412	0.415	—	—	—	—
10^7	10^6	0.0301	0.0283	0.0256	0.177	0.0189	0.00506	0.690	0.0231
10^7	3×10^6	0.434	0.359	0.282	0.448	0.0182	0.0428	0.643	0.180
10^7	10^7	12.1	12.0	62.2	61.5	—	—	—	—

Table 2.2: Median errors in intrinsic quantities for 10^4 binaries of various masses at $z = 3$.

m_1 ($M_\odot$)	m_2 ($M_\odot$)	$\Delta m_1/m_1$	$\Delta m_2/m_2$	$\Delta\chi_1$	$\Delta\chi_2$	$\Delta\mathcal{M}/\mathcal{M}$ (no precession)	$\Delta\mathcal{M}/\mathcal{M}$ (precession)	$\Delta\mu/\mu$ (no precession)	$\Delta\mu/\mu$ (precession)
10^5	10^5	0.00791	0.00792	0.0392	0.0389	—	—	—	—
3×10^5	10^5	0.00811	0.00658	0.0193	0.0359	0.00103	8.00×10^{-5}	0.172	0.00290
3×10^5	3×10^5	0.0134	0.0134	0.0615	0.0616	—	—	—	—
10^6	10^5	0.00718	0.00502	0.00993	0.0409	0.00326	0.000184	0.305	0.00391
10^6	3×10^5	0.0156	0.0124	0.0249	0.0460	0.00427	0.000289	0.469	0.00596
10^6	10^6	0.0424	0.0423	0.197	0.200	—	—	—	—
3×10^6	3×10^5	0.0161	0.0117	0.0158	0.0808	0.0115	0.00103	0.643	0.00922
3×10^6	10^6	0.0576	0.0475	0.0606	0.107	0.0108	0.00265	0.635	0.0214
3×10^6	3×10^6	0.396	0.391	2.43	2.44	—	—	—	—
10^7	10^6	0.279	0.282	0.208	1.41	0.0374	0.0640	0.704	0.232
10^7	3×10^6	10.1	8.41	6.10	7.61	0.106	1.11	0.769	4.28
10^7	10^7	2280	2290	10300	9900	—	—	—	—

Table 2.3: Median errors in intrinsic quantities for 10^4 binaries of various masses at $z = 5$. The results for the highest masses are meaningless — the parameters are completely undetermined.

a dramatic effect. So much of their signal is moved out of band that LISA may measure their waves for only a very short time. As such, measurement may not provide sufficient information to constrain 15 parameters. This is reflected in the high condition numbers associated with such cases. Their Fisher matrices are thus nearly singular, and their inverses are untrustworthy. In fact, measurement error in these binaries actually *degrades* when precession is included. The time in band is too short for precession effects to accumulate. They do not aid parameter estimation; instead, the need to fit extra parameters causes errors to be worse.

2.4.3 Sky position and distance to source

We now focus on extrinsic parameters, the sky position and the luminosity distance to the source. We find that the determination of these parameters is likewise improved when precession physics is taken into account, though not as strongly as for intrinsic parameters. This might be expected, since precession is an intrinsic effect local to the binary and has no direct dependence on these extrinsic parameters. Precession's impact on the extrinsic parameters is somewhat more indirect — it largely improves their determination by reducing the (otherwise quite strong) correlation between sky position and the orbital angular momentum direction $\hat{\mathbf{L}}$ and between these angles and the source's luminosity distance.

In our analysis, a binary's position on the sky is characterized by the two parameters $\boldsymbol{\theta} = (\bar{\mu}_N = \cos\bar{\theta}_N, \bar{\phi}_N)$. We want to convert from errors in these two parameters to an error ellipse on the sky. To do so, we first perform a change of coordinates from $\bar{\mu}_N$ to $\bar{\theta}_N$. For small deviations from the ML estimate, we have $\delta\bar{\theta}_N = (d\bar{\theta}_N/d\bar{\mu}_N)\delta\bar{\mu}_N = -\delta\bar{\mu}_N/\sin\bar{\theta}_N$. Next, we recognize that due to the geometric properties of the sphere, the same $\delta\bar{\phi}_N$ corresponds to a different "proper" angle depending on the value of $\bar{\theta}_N$: $\delta\bar{\phi}^p_N = \sin\bar{\theta}_N\delta\bar{\phi}_N$. With these modifications, equation (2.57) becomes

$$p(\tilde{\boldsymbol{\theta}}|s) \propto \exp\left(-\frac{1}{2}\Gamma^p_{a'b'}\delta\theta^{a'}_p\delta\theta^{b'}_p\right) . \tag{2.74}$$

Here $\delta\theta^{a'}_p \equiv (\delta\bar{\theta}_N, \delta\bar{\phi}^p_N)$ denotes the proper errors accounting for the metric of the

sphere, and $\Gamma^p_{a'b'}$ represents the equivalent Fisher matrix with all conversion factors absorbed inside:

$$\Gamma^p_{\bar{\theta}_N\bar{\theta}_N} = \sin^2\bar{\theta}_N\Gamma_{\bar{\mu}_N\bar{\mu}_N}\,, \tag{2.75}$$

$$\Gamma^p_{\bar{\phi}_N\bar{\phi}_N} = \csc^2\bar{\theta}_N\Gamma_{\bar{\phi}_N\bar{\phi}_N}\,, \tag{2.76}$$

$$\Gamma^p_{\bar{\theta}_N\bar{\phi}_N} = \Gamma^p_{\bar{\phi}_N\bar{\theta}_N} = (-\sin\bar{\theta}_N)(\csc\bar{\theta}_N)\Gamma_{\bar{\mu}_N\bar{\phi}_N} = -\Gamma_{\bar{\mu}_N\bar{\phi}_N}\,, \tag{2.77}$$

and so on for the rest of the elements. The inverse of this matrix is the proper covariance matrix, $\Sigma_p^{a'b'}$. Consider now just the 2×2 subspace of the covariance matrix containing the sky position variables. Let the eigenvalues of this subspace be $\lambda_\pm$. If we define the error ellipse such that the probability that the source lies outside of it is e^{-1} (corresponding to a $\approx 63\%$ confidence interval), then the major and minor axes are given by $2a = 2(2\lambda_+)^{1/2}$ and $2b = 2(2\lambda_-)^{1/2}$, respectively. Expressed in terms of the original covariance matrix, these are

$$\left\{\begin{array}{c} 2a \\ 2b \end{array}\right\} = 2\Bigg[\csc^2\bar{\theta}_N\Sigma^{\bar{\mu}_N\bar{\mu}_N} + \sin^2\bar{\theta}_N\Sigma^{\bar{\phi}_N\bar{\phi}_N} \pm\sqrt{(\csc^2\bar{\theta}_N\Sigma^{\bar{\mu}_N\bar{\mu}_N} - \sin^2\bar{\theta}_N\Sigma^{\bar{\phi}_N\bar{\phi}_N})^2 + 4(\Sigma^{\bar{\mu}_N\bar{\phi}_N})^2}\Bigg]^{1/2}\,. \tag{2.78}$$

We also find the area of the error ellipse:

$$\Delta\Omega_N = \pi ab = 2\pi\sqrt{\Sigma^{\bar{\mu}_N\bar{\mu}_N}\Sigma^{\bar{\phi}_N\bar{\phi}_N} - (\Sigma^{\bar{\mu}_N\bar{\phi}_N})^2}\,. \tag{2.79}$$

Many previous analyses have reported $\Delta\Omega_N$ or $(\Delta\Omega_N)^{1/2}$, the side of a square of equivalent area, as the sky position error [62, 250, 29, 123]. Information about the ellipse's shape, crucial input to coordinating GW observations with telescopes, is not included in such a measure. By examining both $2a$ and $2b$, this information is restored.

Figure 2-7 shows the major axis of the error ellipse $2a$ for both the original code, with no precession, and the code including precession effects. Figure 2-8 shows the

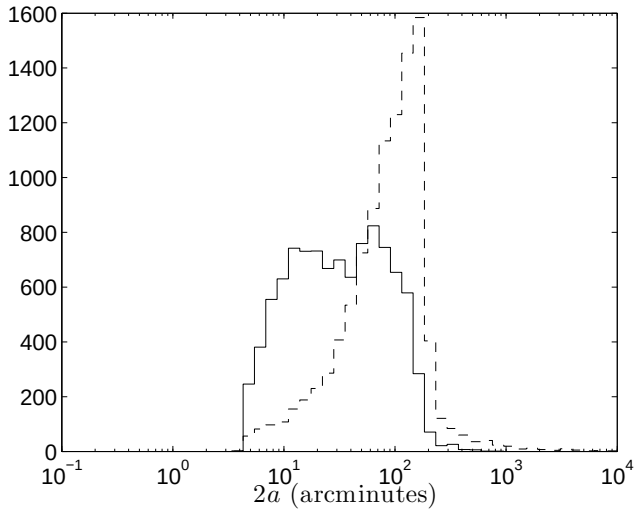

Figure 2-7: Distribution of the major axis of the sky position error ellipse, $2a$, for 10^4 binaries with $m_1 = 10^6 M_\odot$ and $m_2 = 3 \times 10^5 M_\odot$ at $z = 1$. The dashed line is the precession-free calculation; the solid line includes precession. Sky position, as an extrinsic parameter, is improved somewhat indirectly by precession; therefore, the improvement is less than for the masses.

same for the minor axis $2b$. (Note that these figures cannot tell us which major axis is associated with which minor axis; that information is lost in the construction of the histograms.) When precession is included, the median of both distributions is reduced by about half an order of magnitude. The minor axis distribution also shows a long tail of very small errors. In those cases, the position would be very well-constrained in one direction.

Finally, we examine how well distance to the binary is determined. Figure 2-9 compares $\Delta D_L/D_L$ both with and without precession physics taken into account. For this case, the distance error improves by about a factor of 3.

Tables 2.4 and 2.5 show the median extrinsic errors for binaries of different mass at $z = 1$. For comparison purposes, we include results that neglect spin precession. Binaries with the best determined parameters at this redshift have total mass several$\times$ $10^5\, M_\odot \lesssim M \lesssim$ several $\times\, 10^6\, M_\odot$ — smaller binaries are not quite determined so well due to the weakness of their signal, while larger ones are not determined so well because they radiate fewer cycles in band. We also see again that unequal-mass

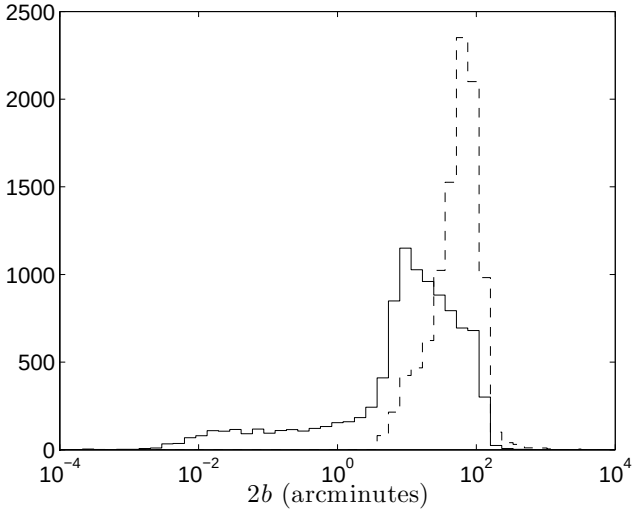

Figure 2-8: Distribution of the minor axis of the sky position error ellipse, $2b$, for 10^4 binaries with $m_1 = 10^6 M_\odot$ and $m_2 = 3 \times 10^5 M_\odot$ at $z = 1$. The dashed line is the precession-free calculation; the solid line includes precession.

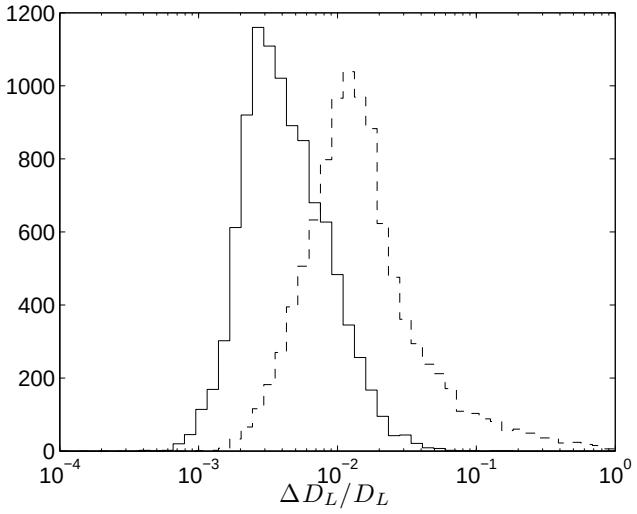

Figure 2-9: Distribution of errors in the luminosity distance for 10^4 binaries with $m_1 = 10^6 M_\odot$ and $m_2 = 3 \times 10^5 M_\odot$ at $z = 1$. The dashed line is the precession-free calculation; the solid line includes precession.

binaries give better results than equal-mass binaries due to the impact of mass ratio on precession effects. Overall, we find that the major axis of the error ellipse is on the order of a few $\times$ 10 arcminutes, while the minor axis is a factor of $2-4$ smaller. This represents an improvement over the "no precession" case by a factor $\sim 2-7$ for the major axis and a factor $\sim 2-10$ for the minor axis. The distance errors are on the order of $0.2\% - 0.7\%$ for most masses, a factor of $\sim 2-7$ improvement.

Tables 2.6–2.9 show the same results for higher redshift. We see the same trends as at $z = 1$, but with some degradation in numerical value. The sky position errors reach a few degrees in the major axis and several tens of arcminutes up to a degree or two in the minor axis. The distance errors are on the order of 1 to several percent for most masses. At the highest masses, we again see that these parameters are essentially undetermined and that precession makes things worse by requiring extra parameters to be fit.

2.5 Summary and conclusions

The general relativistic precession of black holes in a binary system can have a strong influence on the binary's dynamics [10, 139, 214] and thus upon the GWs that it generates. It has been known for some time that it will be necessary to take these dynamics into account in order to detect these black holes in noisy detector data [43, 187, 42, 41, 106, 105, 104] Clearly, taking these dynamics into account will be just as (if not more) important for the complementary problem of determining the parameters which characterize a detected system. Vecchio [250] first demonstrated that, by taking into account precession physics, quite a few near degeneracies among binary source parameters can be broken, making our estimates for how accurately they can be determined more optimistic. This analysis largely confirms and extends Vecchio's pioneering work. By taking the equations of motion to higher order to include spin-spin couplings, and by surveying measurement accuracy as a function of mass ratio, we have found that the improvement noted by Vecchio holds rather broadly. The degeneracy breaking due to precession physics is a rather robust phenomenon.

m_1 ($M_\odot$)	m_2 ($M_\odot$)	$2a$ (arcmin) (no precession)	$2a$ (arcmin) (precession)	$2b$ (arcmin) (no precession)	$2b$ (arcmin) (precession)	$\Delta\Omega_N(\mathrm{deg}^2)$ (no precession)	$\Delta\Omega_N(\mathrm{deg}^2)$ (precession)
10^5	10^5	133	27.3	84.7	13.3	2.44	0.0729
3×10^5	10^5	115	16.9	72.6	7.33	1.81	0.0233
3×10^5	3×10^5	101	23.3	62.8	11.8	1.36	0.0556
10^6	10^5	105	27.2	65.1	6.62	1.47	0.0235
10^6	3×10^5	93.1	31.3	57.5	13.2	1.15	0.0705
10^6	10^6	90.1	40.2	54.1	21.9	1.04	0.176
3×10^6	3×10^5	95.0	34.1	57.3	9.20	1.16	0.0445
3×10^6	10^6	102	32.3	56.0	14.7	1.24	0.0839
3×10^6	3×10^6	135	43.3	68.5	22.3	2.00	0.193
10^7	10^6	149	37.6	75.2	12.2	2.42	0.0670
10^7	3×10^6	238	42.1	119	19.0	6.07	0.142
10^7	10^7	466	81.3	232	38.6	23.3	0.680

Table 2.4: Median errors in sky position for 10^4 binaries of various masses at $z = 1$, including comparisons with the "no precession" case. Note that the given major axis, minor axis, and ellipse area are the medians for each data set and do not correspond to the same binary. However, they still represent an average sky position error ellipse in the following sense: $\sqrt{\pi ab}$, calculated using the median values of $2a$ and $2b$, differs in most cases by less than 10% from the median value of $\sqrt{\Delta\Omega_N}$ (except at more extreme mass ratios — when $m_1/m_2 = 10$, the difference can be 25%).

m_1 ($M_\odot$)	m_2 ($M_\odot$)	$\Delta D_L/D_L$ (no precession)	$\Delta D_L/D_L$ (precession)
10^5	10^5	0.0193	0.00398
3×10^5	10^5	0.0165	0.00240
3×10^5	3×10^5	0.0143	0.00357
10^6	10^5	0.0149	0.00320
10^6	3×10^5	0.0132	0.00393
10^6	10^6	0.0125	0.00560
3×10^6	3×10^5	0.0135	0.00376
3×10^6	10^6	0.0135	0.00419
3×10^6	3×10^6	0.0182	0.00689
10^7	10^6	0.0200	0.00457
10^7	3×10^6	0.0322	0.00610
10^7	10^7	0.0636	0.0250

Table 2.5: Median errors in luminosity distance for 10^4 binaries of various masses at $z = 1$, including comparisons with the "no precession" case.

m_1 ($M_\odot$)	m_2 ($M_\odot$)	$2a$ (arcmin) (no precession)	$2a$ (arcmin) (precession)	$2b$ (arcmin) (no precession)	$2b$ (arcmin) (precession)	$\Delta\Omega_N(\mathrm{deg}^2)$ (no precession)	$\Delta\Omega_N(\mathrm{deg}^2)$ (precession)
10^5	10^5	432	81.0	271	40.8	25.3	0.665
3×10^5	10^5	389	92.5	242	39.5	20.4	0.656
3×10^5	3×10^5	356	142	220	75.7	16.9	2.15
10^6	10^5	379	141	233	36.6	19.0	0.739
10^6	3×10^5	359	129	215	56.7	16.5	1.25
10^6	10^6	416	158	224	84.3	20.4	2.64
3×10^6	3×10^5	425	132	233	40.3	21.8	0.751
3×10^6	10^6	599	142	302	64.6	38.8	1.65
3×10^6	3×10^6	990	224	494	111	106	5.08
10^7	10^6	1320	206	648	78.5	184	2.74
10^7	3×10^6	2380	297	1180	152	621	9.40
10^7	10^7	6820	2000	3390	583	5070	256

Table 2.6: Median errors in sky position for 10^4 binaries of various masses at $z = 3$.

m_1 ($M_\odot$)	m_2 ($M_\odot$)	$\Delta D_L/D_L$ (no precession)	$\Delta D_L/D_L$ (precession)
10^5	10^5	0.0617	0.0123
3×10^5	10^5	0.0551	0.0126
3×10^5	3×10^5	0.0502	0.0201
10^6	10^5	0.0550	0.0155
10^6	3×10^5	0.0500	0.0161
10^6	10^6	0.0556	0.0237
3×10^6	3×10^5	0.0568	0.0153
3×10^6	10^6	0.0809	0.0193
3×10^6	3×10^6	0.134	0.0422
10^7	10^6	0.178	0.0293
10^7	3×10^6	0.326	0.0805
10^7	10^7	0.935	2.41

Table 2.7: Median errors in luminosity distance for 10^4 binaries of various masses at $z = 3$.

m_1 ($M_\odot$)	m_2 ($M_\odot$)	$2a$ (arcmin) (no precession)	$2a$ (arcmin) (precession)	$2b$ (arcmin) (no precession)	$2b$ (arcmin) (precession)	$\Delta\Omega_N(\mathrm{deg}^2)$ (no precession)	$\Delta\Omega_N(\mathrm{deg}^2)$ (precession)
10^5	10^5	729	169	456	85.7	71.8	2.93
3×10^5	10^5	676	217	419	95.8	61.0	3.73
3×10^5	3×10^5	650	295	395	161	54.9	9.29
10^6	10^5	686	248	416	66.8	61.0	2.35
10^6	3×10^5	716	233	404	101	63.0	3.96
10^6	10^6	976	315	497	162	105	10.2
3×10^6	3×10^5	986	265	507	86.4	108	3.27
3×10^6	10^6	1620	304	810	139	282	7.52
3×10^6	3×10^6	2930	538	1460	260	928	29.5
10^7	10^6	5080	577	2480	290	2760	31.9
10^7	3×10^6	10500	1720	5130	621	11900	234
10^7	10^7	75500	180000	35000	29600	618000	1.15×10^6

Table 2.8: Median errors in sky position for 10^4 binaries of various masses at $z = 5$. Again, the results for the highest masses are essentially meaningless—the parameters are completely undetermined.

m_1 ($M_\odot$)	m_2 ($M_\odot$)	$\Delta D_L/D_L$ (no precession)	$\Delta D_L/D_L$ (precession)
10^5	10^5	0.104	0.0260
3×10^5	10^5	0.0957	0.0284
3×10^5	3×10^5	0.0917	0.0409
10^6	10^5	0.0983	0.0273
10^6	3×10^5	0.0961	0.0294
10^6	10^6	0.132	0.0501
3×10^6	3×10^5	0.133	0.0318
3×10^6	10^6	0.220	0.0436
3×10^6	3×10^6	0.400	0.140
10^7	10^6	0.689	0.124
10^7	3×10^6	1.42	1.24
10^7	10^7	10.3	377

Table 2.9: Median errors in luminosity distance for 10^4 binaries of various masses at $z = 5$.

Two conclusions from this work are particularly important with regard to the astrophysical reach of future LISA measurements. The first is that modeling spin-precession physics makes it possible to determine the magnitudes of the spins of the black holes which constitute the binary. If the spins are rapid, they can be measured quite accurately (as good as 0.1% accuracy for high spin, low redshift systems) due to the strong modulation imposed on the signal by their interaction. Coupled with the fact that the black hole masses can likewise be measured with good precision, this suggests that LISA will be a valuable tool for tracking the evolution of both mass and spin over cosmic time. Such observations could provide a direct window into the growth of cosmological structures. Measuring spin may also make it possible to indirectly test the black hole area theorem [111]. The requirement that black hole area can only grow implies a consistency relation between the initial and final masses and spins. By measuring the initial masses and spins through the inspiral, and the mass and spin of the merged remnant hole through the ringdown waves [71, 30], we can check this consistency relation in a manner analogous to the mass loss test proposed in [132]. We intend to investigate whether this test is feasible in future work.

Second, we confirm Vecchio's result that precession breaks degeneracies between the angles which determine a binary's orientation and its position on the sky, improving the accuracy with which sky position can be fixed using GWs alone. At low redshift ($z \sim 1$), we find that sources can be localized to within an ellipse whose major axis is typically a few $\times 10$ arcminutes across and whose minor axis is typically a factor $\sim 2-4$ smaller. This ellipse is small enough that searching it for an electromagnetic counterpart to the coalescence event should not be too arduous a task [142, 144, 157, 143]. For coalescences at higher redshift, the waves weaken and the source is not so well localized. The field which would need to be searched for sources at $z \sim 3-5$ is typically a few degrees across in the long axis and tens of arcminutes to a degree or two in the short direction — a rather more difficult challenge, but not hopeless. In Chapter 3, we investigate the nature of localization with spin precession more thoroughly, including how the errors evolve with observation time up to final merger.

As mentioned in Sec. 2.1, our analysis makes many assumptions and approximations which are likely to affect our results. Because of these simplifications, we cannot claim that this analysis gives a definitive statement about the accuracy with which LISA could measure binary black hole source parameters. However, it is certainly *indicative* of the accuracy which we expect LISA to achieve. In particular, we are confident that the trends we have seen as parameters are varied (e.g., masses, redshift, spin magnitude) are robust. Most importantly, it is very clear that the influence of spin-induced precession upon the measured waveform allows parameters to be measured to greater accuracy than before.

A goal of future work will be to lift the approximations. One major concern is the Gaussian approximation we have taken to the likelihood function. As already discussed, this approximation is known to be good when the SNR is "large" [90, 63]; however, it is not apparent what large really means, particularly given that we are fitting for 15 parameters. Lifting this simplifying approximation can be done by simply computing the likelihood function (2.55) directly and examining how well parameters are thereby determined. In the context of GW measurements, Markov chain Monte Carlo (MCMC) techniques have been investigated and found to be very useful [52, 51, 247]; the application of these techniques to LISA measurement problems is now being rather actively investigated [58, 59, 60]. With Neil Cornish and Samaya Nissanke, we have recently developed an MCMC code including precession physics. Preliminary results indicate that the Gaussian approximation works well for most parameters; however, it will be interesting to investigate the dependence on spin and SNR.

Using this code as a basis, Cornish has also developed a newer Fisher code[14] which contains both precession physics and higher harmonics [15]. Both of these effects have now been shown to improve parameter estimation errors by a significant amount. Since these two improvements arise from very different physical effects, it is likely that their separate improvements can be combined for an overall improvement

[14]Incidentally, both of these codes also abandon the stationary phase approximation, while incorporating a more realistic LISA response and noise curve. While we are not testing the effect of these changes directly, they certainly make the end results more realistic.

significantly better than each effect on its own. Unfortunately, preliminary results show the opposite: Once a degeneracy has been broken, it cannot be broken any further. Still, we have only barely begun to explore this problem, and differences between the effects of precession and higher harmonics may yet emerge.

Finally, another new version of our code developed by Stephen O'Sullivan has shown that including higher-order phase and precession terms in the waveform model has relatively little impact on parameter errors. For small mass ratios, including these higher-order terms actually degrades the errors by ~ 20 %, while for $m_1/m_2 \gtrsim 3$, the errors are improved by ~ 10 % [185]. It seems that the lowest-order spin effects described in this chapter contribute the most to LISA's parameter estimation capability.

Chapter 3

Localization of massive black hole binaries

3.1 Introduction

3.1.1 The LISA GW pixel

Among the most important sources of gravitational waves (GWs) in the low-frequency band of space-based detectors are the coalescences of massive black hole binaries (MBHBs). As we saw in the previous chapter, "intrinsic" parameters — the masses and spins of the black holes which compose the binary — should be determined with very high accuracy, with relative errors typically $\sim 10^{-3}$ to 10^{-1}, depending on system mass and redshift. By measuring an ensemble of coalescences over a range of redshifts, MBHB GWs may serve as a kind of structure tracer, tracking the growth and spin evolution of black holes over cosmic time.

"Extrinsic" system parameters, describing a binary's location and orientation relative to the detector, are also determined by measuring its GWs. In the previous chapter, we showed that a binary's position on the sky can be localized at $z = 1$ to an ellipse with a major axis of a few tens of arcminutes and a minor axis a factor of $2 - 4$ smaller. At higher redshift ($z = 3 - 5$), these values degrade by a factor of a few, reaching several degrees in the long direction and tens of arcminutes to a degree

or two in the short one. We also found that a source's luminosity distance typically can be determined to better than 1% at low redshift ($z = 1$), degrading to several percent at higher redshift ($z = 3 - 5$).

The intrinsic ability of GWs to determine the distance to a coalescing binary is phenomenal. Coalescing MBHB systems constitute exquisitely well-calibrated distance measures, with the calibration provided by general relativity. Unfortunately, in practice this percent-level or better accuracy could only be achieved if we measured MBHB coalescences in an empty universe. In our universe, weak lensing will magnify or demagnify the GWs, and we will infer a luminosity distance smaller (for magnification) or larger (for demagnification) than the true value. This phenomenon affects all high-redshift standard candles. Its impact on Type Ia supernovae in particular has been discussed in detail [98, 125, 121, 124].

The distance error scales with redshift roughly as $(\Delta D_L/D_L)_{\rm lens} \simeq 0.044z$ for low z [124]. It is expected that this dependence will become flat at some transition redshift, most likely near $z \sim 3$ or 4; the precise transition depends upon the (poorly understood) high-redshift mass function [122]. With the development of high-quality weak lensing maps, one might think that it would be possible to correct for the impact of lensing and recover much of the intrinsic GW distance measurement precision. Unfortunately, lensing noise arises mostly from structure on sub-arcminute scales that is not probed by shear maps, making any substantial correction impossible [64]. Since we will not know the extent of the magnification when we measure MBHB waves, we must simply accept the fact that lensing introduces a dispersion of several percent in determining the distance to these GW events. Practically, we can compute

$$\frac{\Delta D_L}{D_L} \simeq \sqrt{\left(\frac{\Delta D_L}{D_L}\right)^2_{\rm GW} + \left(\frac{\Delta D_L}{D_L}\right)^2_{\rm lens}} \simeq \max\left[\left(\frac{\Delta D_L}{D_L}\right)_{\rm GW}, \left(\frac{\Delta D_L}{D_L}\right)_{\rm lens}\right] . \quad (3.1)$$

When we quoted distance measurement errors in Chapter 2, we only quoted the intrinsic GW measurement error $(\Delta D_L/D_L)_{\rm GW}$, neglecting lensing's impact. We continue that practice in this chapter. However, we note that whenever the intrinsic GW distance error is $\lesssim 5\%$, lensing will blur it to the several percent level.

Note that a source's redshift z *cannot* be directly determined using only GWs. Gravitational wave measurements infer system parameters through their impact on certain dynamical time scales, such as orbital frequencies and the rate at which these frequencies evolve. Since these time scales all suffer cosmological redshift, z is degenerate with other parameters. For example, any mass parameter m is actually measured as $(1+z)m$ (the "redshifted mass"). However, if the binary's luminosity distance is determined, its redshift can then be inferred by assuming a cosmography. If cosmological parameter errors can be neglected, then one typically finds that the redshift error is about equal to the distance error [127]: $\Delta z/z \approx \Delta D_L/D_L$, independent of redshift. For most binaries, the redshift can be determined to several percent (with an error budget dominated by gravitational lensing[1]). We thus expect that GW measurements will locate a binary to within a three-dimensional "GW pixel" which at $z = 1$ has a cross-sectional area of $\sim 10^{-2}$ to 10^{-1} deg^2 and a depth $\Delta z/z \sim$ several percent.

3.1.2 Electromagnetic counterparts to MBHB GW events

It is anticipated that there will be great interest in searching the GW pixel for electromagnetic (e.g., optical, X-ray, radio) counterparts to MBHB GW events. It is plausible that no significant electromagnetic activity occurs in conjunction with MBH coalescence. In this case, one could imagine searching the pixel for a galaxy with a structure that indicates a recent merger. Another possibility is to search for a galaxy which has a bulge radial velocity consistent with the GW-measured final black hole mass. (This of course assumes that the well-known relation between these quantities in the local universe [87, 99] holds at high redshift and so soon after a merger.) Given the typical size of the LISA pixel, such searches will be quite difficult.

It is likely, however, that there is some unique EM activity associated with the MBH coalescence. The nature of such activity has become a hot research topic in

[1]At redshifts $z \lesssim 0.3$, the error is actually dominated by peculiar velocity effects [142]; however, the event rate is probably negligible at such low redshifts. As such, we will focus on gravitational lensing as the main source of systematic redshift error.

the past few years, leading to many possible scenarios [70, 143]. For example, if the surrounding gas is completely swept away by the binary, there may be no signal during the coalescence itself. Instead, there would be a delayed afterglow when the gas later accretes onto the remnant hole [177]. This afterglow may occur years after the merger. It is likely, though, that the gas will not be totally swept away, leaving enough to accrete onto the holes and create variable EM activity during the inspiral phase. For example, Armitage & Natarajan showed that for a large-mass-ratio binary, any gas which does remain will be driven in during inspiral, producing an EM signal [11]. More recent work by MacFadyen & Milosavljević showed that periodic variations in the Newtonian potential can create a quasi-periodic EM flux [167].

Other scenarios predict transient signals during or immediately after the merger. Recent work by Bode and Phinney [38] suggests that the final burst of radiation from a coalescing binary (which can convert ~ 10 % of the system's mass to GWs very suddenly) may excite radial waves, and consequently electromagnetic variability, in an accretion disk due to the quick change in the disk's Keplerian potential. Another possibility is that the "kick," or momentum imparted to the black hole due to an asymmetric emission of GWs [92, 46, 215], will send the remnant through the surrounding gas, producing shocks [165, 227, 216]. A transient signal might also appear when the GWs are viscously dissipated in the surrounding gas [145].

Finding a counterpart could greatly enhance the science return of MBH measurement. For example, counterparts can improve LISA's ability to determine certain parameter values. As we have seen in Chapter 2, the sky position is strongly correlated with various other parameters, particularly luminosity distance and orbit orientation. When it is determined exactly by identification of a counterpart, the other parameters can be estimated to greater accuracy [127, 123]. Another difficulty with parameter estimation is that the estimated masses are redshifted masses, not rest-frame masses. The GWs themselves only give the luminosity distance, not the redshift, so any decoupling of mass and redshift will be dependent on a cosmological model [127]. The counterpart gives the redshift directly.

Finding the redshift from a counterpart may also allow us to use MBH events

as cosmological distance measures. Combining the EM-measured redshift with the GW-measured luminosity distance creates a Hubble diagram which is calibrated only by general relativity [217, 123]. In reality, the precision of such "standard sirens" is degraded by weak lensing uncertainties. Still, the systematics affecting MBH GWs should be completely different from those affecting Type Ia supernova standard candles and could serve as a useful complement to those sources.

Counterparts are also useful for studying the astrophysics of the MBH coalescence; indeed, the sheer variety of counterpart scenarios shows how uncertain these processes are. Specifically, counterparts may give insight into gas dynamics and accretion. For instance, GW measurements of the mass and EM measurements of the luminosity can be combined to find the Eddington ratio, $L/L_{\rm Edd}$ [142]. Finally, the counterparts could be used to test fundamental physics. If a counterpart features EM variation in phase with the gravitational wave signal, the two signals can be compared to test the equivalence of photon and graviton propagation speed. Any difference could be explained by a nonzero graviton rest mass [143].

3.1.3 Outline of this chapter

The purpose of this chapter is to examine the localization of MBHB systems more thoroughly, in particular how the GW pixel evolves as the final merger is approached. Chapter 2 only presented results for measurements that proceed all the way to merger. It will clearly be of some interest to monitor potential hosts for the binary event some time before the merger happens; if nothing else, telescopes will need prior warning to schedule observing campaigns. Understanding the rate at which localization evolves can also have an important impact on the design of the LISA mission, clarifying how often it will be necessary to downlink data about MBHB systems in order to effectively guide surveys.

Our main goal is to understand for what range of masses and redshifts prior localization of a binary using GWs will be possible. A previous analysis by Kocsis et al. [144] (hereafter KHMF) also examined this problem in great detail, but without including the impact of spin-induced precession. One of our goals is to see to what

extent precession physics changes the conclusions of KHMF. We find that precession has a fairly small impact on the time evolution of the GW pixel except in the last few days before the final merger, at which point its impact can be tremendous. Precession typically changes the area of the sky position error ellipse by a factor of $\sim 3-10$ (up to ~ 60 in extreme cases) in just the final day. This is in accord with the predictions of KHMF (and even earlier predictions by Neil Cornish [56]).

The necessary background for this chapter can be found in the previous one. We briefly review the localization results of that chapter in Sec. 3.2. The "precession" results of Tables 2.4–2.9 are discussed in more detail, and the distribution of errors at merger is investigated for different mass ratios.

Our main results are presented in Sec. 3.3. We begin by summarizing the key ideas behind the "harmonic mode decomposition" of KHMF in Sec. 3.3.1. This technique cleverly allows calculation of the GW pixel and its time evolution with much less computational effort than our method (albeit without including the impact of spin precession). Unfortunately, we have discovered that some of the approximations used by KHMF introduce a systematic underestimate of the final sky position error by a factor of $2-4$ or more in angle; the approximations are much more reliable a week or more prior to the black holes' final merger. Modulo this underestimate, the KHMF results agree well with our "no precession" code (particularly a week or more in advance of merger, when their underestimate is not severe). KHMF thus serves as a useful point of comparison to establish the impact of precession on source localization.

Section 3.3.2 is dedicated to our own time-dependent localization results, including comparison to KHMF when appropriate. We find that all relevant parameter errors decrease slowly with time until the last day before merger, when they drop more dramatically. This sudden drop is not found in KHMF, nor is it present in a variant of our analysis that ignores spin precession. It clearly can be attributed to the impact of precession on the waveform. Before this last day, the major axis is $\sim 1.5-6$ times, the minor axis $\sim 2-9$ times, and the intrinsic error in the luminosity distance $D_L \sim 1.5-7$ times bigger than at merger for most binaries (i.e., all except the highest masses). Going back to one week (one month) before merger, these numbers

change to $2-9$ $(4-11)$ for the major axis, $3-14$ $(5-24)$ for the minor axis, and $3-14$ $(5-18)$ for the error in the luminosity distance. As a result, for $z=1$, most binaries can be located within a few square degrees a week before merger and 10 deg^2 a month before merger. The intrinsic distance errors are also small enough this early that $\Delta z/z$ remains dominated by gravitational lensing errors of several percent. Advanced localization of MBHB coalescences thus seems plausible for these binaries; the situation is less promising for sources at higher redshift.

As a corollary to our study of the time evolution, we also examine the sky position dependence of errors in Sec. 3.3.3. The errors depend strongly on the polar angle with respect to the ecliptic, increasing in the ecliptic plane to as much as 35% over the median for the major axis, 85% over the median for the minor axis, and 15% over the median for errors in the luminosity distance. The errors have a much weaker dependence on the azimuthal angle. When we convert to Galactic coordinates, we find that the best localization regions appear to lie fairly far out of the Galactic plane, offering hope that searches for counterparts will not be too badly impacted by foreground contamination.

We conclude this chapter in Sec. 3.4. Besides summarizing our results, we discuss shortcomings of this analysis and future work which could help to better understand how well GWs can localize MBHB sources.

3.2 The LISA pixel at merger

We saw in Chapter 2 that LISA's ability to measure the sky position of a source is primarily due to its motion around the Sun. (By comparison, ground-based detectors rely on comparing wave arrival times at different instruments.) This motion creates two effects: a changing orientation of the detector relative to the source and a Doppler shift in the frequency of the waves. However, even with these modulations, it is difficult to detangle sky position from the orientation of the binary. Additional modulations due to spin-induced precession help to break this degeneracy. As we saw in Chapter 2, including precession effects can improve sky position errors by about

half an order of magnitude in each direction. Because the luminosity distance is strongly correlated with the sky position, measurements of D_L improve by about the same factor.

All of the results in Chapter 2 are computed at the end of inspiral. Since sky position is primarily determined by the motion of LISA around the Sun, one might guess that the inspiral provides the bulk of the localization. However, with the development of numerical relativity, it has become possible to perform parameter estimation calculations which include the merger phase of the waveform. Preliminary results seem to show that the merger could dramatically improve sky position and distance determination [18, 242]. Nevertheless, we shall consider our end-of-inspiral results to be the "final" accuracy with which MBHB sky position and distance can be determined with GWs. In this section, we summarize these results in order to give a baseline for the time-dependent results of Sec. 3.3.

LISA's ability to localize a source at merger is summarized in Tables 2.4–2.9 of Chapter 2, specifically the "precession" columns. (In this chapter, all results will include precession physics unless otherwise noted.) We see that for all redshifts, the accuracy is worst for the largest masses. This is because the most massive systems are in band for the least amount of time. The short time these systems spend in band means that LISA measures a relatively small number of modulations (whether induced by the constellation's orbit or by spin precession). Second, note that the results for $m_1 = m_2$ tend to be less accurate than results with similar total mass but for which $m_1 > m_2$. As we saw in Chapter 2, the cause of this phenomenon lies in the precession equations (2.15) and (2.16): When $m_1 \neq m_2$, the two spins precess at different rates, imposing richer modulations on the measured GWs. Since $m_1 = m_2$ is a rather artificial limit, we expect that the more accurate results for nonunity mass ratio will be the rule.

Independent of these trends, an important result is that MBHB systems are pinned down on the sky fairly accurately at $z = 1$. Modulo the higher-mass binaries, the median major axis of the sky position error ellipse is typically about $15' - 45'$, and the median minor axis is about $5' - 20'$, with a total ellipse area considerably smaller

than 1 deg^2 (ranging from about 0.02 to 0.2 deg^2). Sources at $z = 1$ are located accurately enough that one can comfortably contemplate searching the GW error ellipse for MBHB counterparts with future survey instruments, such as the Large Synoptic Survey Telescope (LSST) [246].

At higher redshift, positional accuracy degrades. This is due to the weakening of the signal with distance and to the redshifting of the waves' spectrum, so that the signal tends to spend less time in band. At $z = 3$, the major axis of the error ellipse is $\sim 1^\circ - 4^\circ$ across, and the minor axis is $\sim 40' - 110'$. The total area of this ellipse is $\sim 0.5 - 5$ deg^2. At $z = 5$, this degrades further to $\sim 3^\circ - 5^\circ$ for the major axis, $\sim 1^\circ - 3^\circ$ for the minor axis, and $\sim 2 - 10$ deg^2 for the total area. These degraded accuracies are still sufficiently small that telescopic searches for MBHB counterparts have a good chance to be fruitful (although not nearly as simple as they would be at $z \sim 1$).

In all cases, the GW distance determination is extremely accurate: For all but the highest masses, $\delta D_L/D_L \lesssim 0.7\%$ at $z = 1$, $\lesssim 4\%$ at $z = 3$, and $\lesssim 5\%$ at $z = 5$. Distance is determined so precisely that these errors are in fact irrelevant — weak gravitational lensing will dominate the distance error budget for all but the most massive MBHB events.

Although the median values reported in Tables 2.4–2.9 accurately indicate the typical localization accuracies we expect, it should be emphasized again that they are taken from broad distributions, Figure 3-1 presents the major and minor axis distributions we computed for binaries at $z = 1$ with masses $m_1 = 10^6\, M_\odot$ and $m_2 = (10^5, 3 \times 10^5, 10^6) M_\odot$. (The middle case was also pictured in Figures 2-7 and 2-8, in comparison to the "no precession" results.) Note that the major axis distribution (top) is rather flat when compared to the minor axis distribution (bottom). It lacks a single well-defined peak; in fact, it is actually bimodal for $m_1/m_2 > 1$, with one peak near $10' - 20'$ and another closer to $1^\circ - 2^\circ$. We find that this behavior holds over all mass and redshift cases of interest, with only slight variations. At smaller masses, the distribution is broader than the cases pictured, without strong bimodality; for larger masses, the distribution is somewhat narrower and tends to develop two very

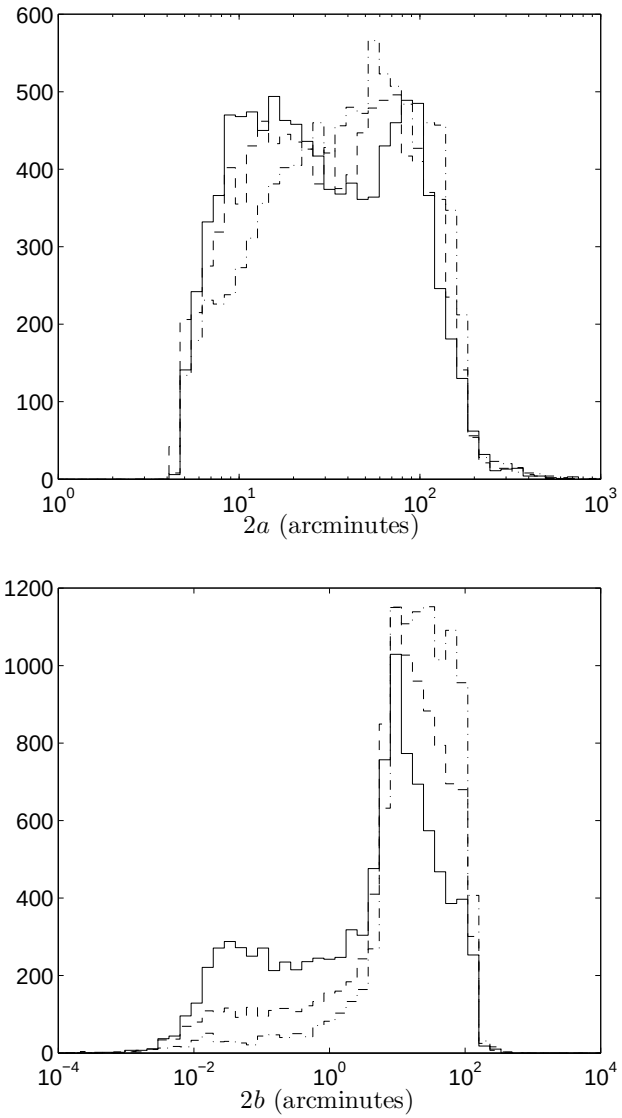

Figure 3-1: Distribution of the major axis $2a$ (top) and minor axis $2b$ (bottom) of the sky position error ellipse for 10^4 binaries with $m_1 = 10^6 M_\odot$ and $m_2 = 10^5 M_\odot$ (solid line), $3 \times 10^5 M_\odot$ (dashed line), and $10^6 M_\odot$ (dash-dotted line) at $z = 1$. Note the bimodal character of the major axis as well as the long tail down to small minor axis, both of which are particularly prominent for larger mass ratios.

distinct peaks. Higher mass ratios tend to accentuate the peaks. These results hold for higher redshift, except that transitions between various behaviors occur at smaller total mass (in keeping with the fact that it is not mass but $(1+z)$ times the mass that determines dynamical behavior).

The minor axis distribution exhibits a rather long tail to very small values, especially when the mass ratio is large. For $m_1/m_2 = 10$, the distribution peaks near a minor axis $\sim 10'$, but extends from roughly $1''$ to about $100'$. As the mass ratio approaches 1, the peak moves to slightly larger values (slightly more than $10'$ for $m_1/m_2 = 3$; roughly $30'$ for $m_1/m_2 = 1$), and the tail becomes less populated (although the distributions span roughly the same extent as when $m_1/m_2 = 10$). We find that this tail exists for all interesting mass and redshift combinations, with the same strong dependence on mass ratio as shown in Figure 3-1.

3.3 Time and position dependence of the LISA pixel

We turn now to a detailed discussion of how well GWs localize an MBHB system as a function of time before final merger and as a function of sky location. We begin by discussing the analysis of KHMF, which presents a clever algorithm for estimating extrinsic parameter errors as a function of time until merger (although at present it does not include spin precession). We demonstrate that their analysis unfortunately underestimates final position errors by roughly a factor of $\sim 2-4$ or more (in angle) due to neglect of certain parameter correlations; the underestimate is much less severe a week or more prior to merger. We then present our own results for the time dependence of the GW pixel, using KHMF for comparison where appropriate. Finally, we conclude with a brief study of the pixel's dependence on sky location.

3.3.1 Summary of KHMF

KHMF have devised a new method, the harmonic mode decomposition (HMD), to solve for the extrinsic parameter errors as a function of time to merger. In the HMD, modulations caused by LISA's motion are decoupled from the much faster inspiral time scales and are then expanded in a Fourier series. The resulting expression for the measured signal features a time-dependent piece with no dependence on the extrinsic parameters and a time-independent piece with all the parameter dependence. As a result, when Monte Carlo simulations are done across parameter space, the time-dependent integrals do not need to be recomputed for each sample of the distribution. This makes it possible to quickly survey the estimated parameter errors across a wide range of parameter space.

As already emphasized, the waveform model used by KHMF does not (yet) include the impact of spin precession. As such, we intend to use their results as a baseline against which the impact of spin precession can be compared. Before doing so, we first checked to make sure that their results were in agreement with a variant of our code which does not include spin precession [127]. To our surprise, we found that the final position accuracy predicted by KHMF was typically a factor ~ 2 (in angle) more accurate than our code predicted.

After detailed study of the HMD algorithm and comparison with our (precession-free) code, we believe we understand the primary source of this disagreement. KHMF define a set of "slow" parameters $\boldsymbol{\theta}_{\rm slow}$, which correspond (with some remappings) to our extrinsic parameters: $\ln D_L$, $\cos\bar{\theta}_L$, $\cos\bar{\theta}_N$, $\bar{\phi}_L$, and $\bar{\phi}_N$. KHMF also define a set of "fast" parameters $\boldsymbol{\theta}_{\rm fast}$, which, modulo the exclusion of spin, map to our intrinsic parameters[2]: t_c, Φ_c, $\ln\mathcal{M}$, and $\ln\eta$. In their formulation of the HMD, KHMF approximate the cross-correlation between intrinsic and extrinsic parameters to be zero. Although the correlations between intrinsic and extrinsic parameters tend to be small, they are not zero. We find that they typically range in magnitude from about 0.1 to 0.4, sometimes reaching ~ 0.8. Neglecting these correlations altogether

[2]In Chapter 2, our precession-free code used $\ln\mu$ as the second mass parameter, but here we switch to $\ln\eta$ in order to match the choice of KHMF.

leads to a systematic underestimate in extrinsic parameter errors.

An example of this is shown in Tables 3.1 and 3.2. To produce the data shown in Table 3.1, we compute the Fisher matrix $\Gamma^{\rm tot}_{ab}$ and then invert for the covariance matrix, $\Sigma^{ab} = (\Gamma^{-1}_{\rm tot})^{ab}$. Table 3.1 then presents a slightly massaged representation of this matrix: Diagonal elements are the 1-σ errors $(\Sigma^{aa})^{1/2}$, and off-diagonal elements are the correlation coefficients $c^{ab} = \Sigma^{ab}(\Sigma^{aa}\Sigma^{bb})^{-1/2}$. Take particular note of the magnitude of the correlations between intrinsic and extrinsic parameters (the upper right-hand portion of Table 3.1). Many entries have values $\sim 0.2 - 0.3$, and two have values $\sim 0.7 - 0.8$.

To see what effect neglecting the intrinsic-extrinsic correlations has, we repeat this exercise, with a slight modification: We compute $\Gamma^{\rm tot}_{ab}$ as before, but we now set to zero entries corresponding to mixed intrinsic/extrinsic parameters. For example, we set by hand $\Gamma^{\rm tot}_{\ln D_L,\ \ln \mathcal{M}} = 0$. We then invert this matrix to obtain Σ^{ab}. The result is shown in Table 3.2. Note that mean parameter error (diagonal entries) is often significantly smaller than errors when these correlations are not ignored. The impact of correlations between intrinsic and extrinsic parameters is clearly not negligible.

Table 3.3 gives further examples illustrating the impact of neglecting these correlations on our estimates of LISA's localization accuracy. We show 10 points drawn from a 10^4 binary Monte Carlo run; all use the same masses and redshifts ($m_1 = 3\times10^6\,M_\odot$, $m_2 = 10^6\,M_\odot$, and $z = 1$) but have different (randomly distributed) sky positions, orientations, spins, and t_c. For these parameters, we find that neglecting intrinsic-extrinsic correlations causes one to underestimate the major axis of the position ellipse by a (median) factor of ~ 2 and the minor axis by a factor of $\sim 3 - 4$; the area is underestimated by a factor of $\sim 6 - 7$.

Our conclusion is that the HMD technique developed by KHMF is overly optimistic by a factor of $\sim 2-4$ or more (in angle) regarding the final accuracy with which GWs can locate an MBHB event on the sky. As a prelude to the time evolution study we present in Sec. 3.3.2, we also examined how this underestimate evolves as merger is approached. To our relief, it appears that this underestimate is *much* smaller prior to merger: For the handful of cases we examined, the factor of $2 - 4$ underestimate

	$\ln D_L$	$\cos\theta_L$	$\cos\theta_N$	$\bar{\phi}_L$	$\bar{\phi}_N$	t_c	Φ_c	$\ln\mathcal{M}$	$\ln\eta$	β	σ
$\ln D_L$	0.233	−0.984	0.878	0.509	0.213	0.0801	0.246	0.227	−0.186	0.205	−0.106
$\cos\bar{\theta}_L$	·	0.467	−0.861	−0.350	−0.071	−0.040	−0.098	−0.178	0.138	−0.156	0.0622
$\cos\bar{\theta}_N$	·	·	0.0006	0.465	0.203	0.0709	0.231	0.201	−0.166	0.181	−0.095
$\bar{\phi}_L$	·	·	·	0.687	0.782	0.232	0.827	0.337	−0.317	0.328	−0.259
$\bar{\phi}_N$	·	·	·	·	0.0017	0.193	0.691	0.252	−0.244	0.250	−0.210
t_c	·	·	·	·	·	63.1	0.705	0.923	−0.955	0.942	−0.993
Φ_c	·	·	·	·	·	·	4.00	0.742	−0.747	0.748	−0.726
$\ln\mathcal{M}$	·	·	·	·	·	·	·	0.0010	−0.995	0.998	−0.956
$\ln\eta$	·	·	·	·	·	·	·	·	0.303	−0.999	0.981
β	·	·	·	·	·	·	·	·	·	1.11	−0.971
σ	·	·	·	·	·	·	·	·	·	·	0.722

Table 3.1: Example of errors (diagonal elements) and correlations (off-diagonal elements) for a binary with $m_1 = 3 \times 10^6\, M_\odot$ and $m_2 = 10^6\, M_\odot$ at $z = 1$. The errors in $\bar{\phi}_L$, $\bar{\phi}_N$, and Φ_c are measured in radians; the error in t_c is measured in seconds. This example was taken from the same Monte Carlo distribution used to make Table 3.3; in this particular case, the randomly distributed parameters have the values $\cos\bar{\theta}_L = -0.628$, $\cos\bar{\theta}_N = 0.850$, $\bar{\phi}_L = 3.50$ rad, $\bar{\phi}_N = 0.514$ rad, $t_c = 6.90 \times 10^7$ s, $\beta = 1.48$, and $\sigma = 0.107$. Entries containing $\cdot$ can be found by symmetry.

	$\ln D_L$	$\cos\theta_L$	$\cos\theta_N$	ϕ_L	ϕ_N	t_c	Φ_c	$\ln\mathcal{M}$	$\ln\eta$	β	σ
$\ln D_L$	0.142	-0.999	0.721	0.999	0.0602	0	0	0	0	0	0
$\cos\bar{\theta}_L$	·	0.294	-0.721	-0.999	-0.0611	0	0	0	0	0	0
$\cos\bar{\theta}_N$	·	·	0.0004	0.719	0.0576	0	0	0	0	0	0
$\bar{\phi}_L$	·	·	·	0.120	0.0664	0	0	0	0	0	0
$\bar{\phi}_N$	·	·	·	·	0.0010	0	0	0	0	0	0
t_c	·	·	·	·	·	61.3	0.990	0.929	-0.959	0.946	-0.993
Φ_c	·	·	·	·	·	·	2.10	0.969	-0.988	0.981	-0.999
$\ln\mathcal{M}$	·	·	·	·	·	·	·	0.0010	-0.995	0.998	-0.960
$\ln\eta$	·	·	·	·	·	·	·	·	0.288	-0.999	0.983
β	·	·	·	·	·	·	·	·	·	1.05	-0.974
σ	·	·	·	·	·	·	·	·	·	·	0.697

Table 3.2: Example of errors and correlations for the same binary as shown in Table 3.1 if correlations between intrinsic and extrinsic parameters are neglected.

$2a$ (arcmin)		$2b$ (arcmin)		$\Delta\Omega_N$ (deg^2)	
Full	KHMF	Full	KHMF	Full	KHMF
201	63.5	59.0	15.1	2.59	0.210
165	120	108	88.6	3.90	2.32
117	61.6	81.0	14.3	2.07	0.193
197	69.7	46.7	13.4	2.01	0.204
10.9	7.23	8.21	5.03	0.0196	0.00793
197	51.2	55.2	13.3	2.37	0.149
46.7	26.8	36.5	9.19	0.372	0.0538
18.1	12.7	15.0	6.37	0.0595	0.0177
155	92.4	88.5	16.2	2.98	0.326
146	143	139	10.9	4.43	0.342

Table 3.3: Ten Monte Carlo points comparing sky position accuracy calculated using the full Fisher matrix technique to that calculated using the KHMF approximation. All ten binaries have $m_1 = 3 \times 10^6\, M_\odot$ and $m_2 = 10^6\, M_\odot$ at $z = 1$.

in angle falls to a mere $10\% - 25\%$ offset one week prior to merger. We find that the offset plateaus at this level, remaining at a few tens of percent up to 28 days before merger.

Accounting for this systematic underestimate, we thus find KHMF's results to be a good baseline against which to compare our results. This comparison makes it possible to assess the extent to which spin precession improves our ability to locate massive black hole binaries prior to the final merger.

3.3.2 Results I: Time evolution of localization accuracy

We finally come to the main results of this chapter, the time evolution of our ability to localize MBHB systems using GWs when spin precession is included. The results summarized in Sec. 3.2 describe the size of the GW pixel (sky position error ellipse and luminosity distance) at the end of inspiral. The end-of-inspiral ("final") localization accuracies are good enough that searching the GW pixel for counterparts to MBHB coalescences is likely to be fruitful. However, given how little is understood about electromagnetic counterparts to these events, it is unclear if waiting until these final moments is the best strategy for such a search. It will surely be desirable to also monitor the best-guess location some days or weeks in advance for electromagnetic

precursors to the final merger. The rate at which the GW pixel evolves as we approach the merger will have strong implications for determining the rate at which LISA data is sent to the ground.

To this end, we now examine the time dependence of the LISA pixel. We have modified the code from Chapter 2 to stop the calculation at a specified time before the fiducial "merge frequency," equation (2.72). We still begin the evolution of each binary at the moment it enters the LISA band (which we take to occur, as before, at $f_{\rm min} = 3 \times 10^{-5}\,{\rm Hz}$). Because we randomly distribute t_c over our (assumed) 3 yr LISA mission, some sources are already in band at the mission's start; consequently, these sources begin at $f > f_{\rm min}$. The binary's evolution is then followed until it reaches a GW frequency $f_{\rm stop} = f(t(f_{\rm merge}) - N)$, where N is the number of days before merger that we want to stop the signal. (Choosing $N = 0$ duplicates the analysis of Chapter 2.)

Figure 3-2 shows the error ellipse evolution for nine examples taken from a sample of 10^4 computed with $m_1 = 10^6 M_\odot$, $m_2 = 3 \times 10^5 M_\odot$ and $z = 1$. We show results for $N = 0$, 1, 2, 4, 7, 14, 21, and 28. For each binary, the major axis is plotted on the x-axis, while the minor axis is plotted on the y-axis; we do not show how each ellipse would be oriented on the sky. (In addition, subsequent error ellipses on the real sky would not be centered at the same point, because the maximum likelihood estimate would change.) The results shown in Figures 3-2a and 3-2b were selected by hand from the distribution as examples of contrasting behavior. The binary in Figure 3-2a shows a dramatic change in the error ellipse with time, especially in the last day before merger. By the end of that day, the binary is localized to an ellipse with $2a = 6.67'$ and $2b = 6.25'$, an area ~ 60 times smaller than at $N = 1$. By contrast, the binary in Figure 3-2b shows almost no change in the error ellipse over the entire four weeks prior to merger.

These are clearly extreme cases. Other extremes exist, including binaries with a minor axis orders of magnitude smaller than the major axis (see the tail in Fig. 3-1, bottom), binaries where the evolution of one or both axes is not strictly monotonic, and binaries which have very large ellipses (essentially filling the sky) for large N.

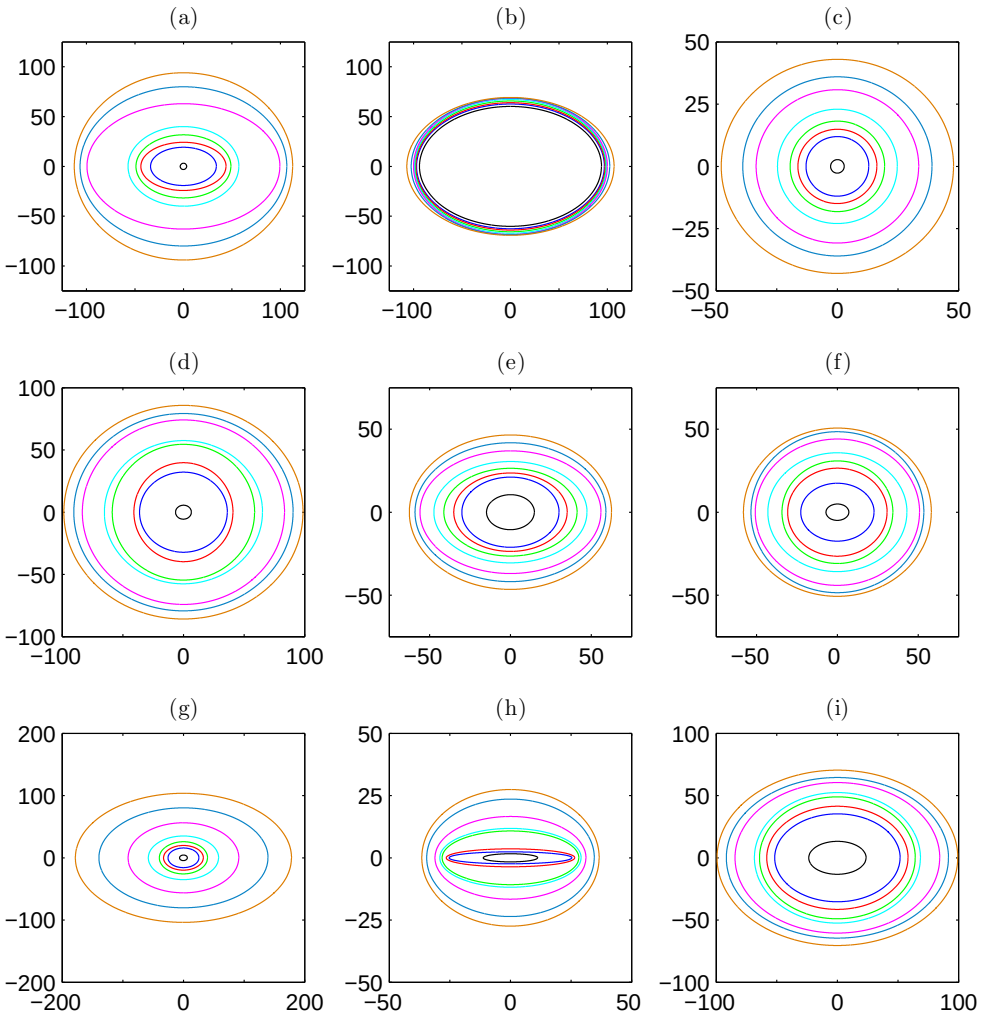

Figure 3-2: Evolution of the sky position error ellipse for nine individual binaries selected from a set of 10^4. All have $m_1 = 10^6 M_\odot$, $m_2 = 3 \times 10^5 M_\odot$, and $z = 1$. The ellipses are oriented so their major axes are parallel to the x-axis and their minor axes are parallel to the y-axis; the axes are labeled in arcminutes. From outside in, the ellipses are evaluated at 28, 21, 14, 7, 4, 2, 1, and 0 days before merger.

(Such cases correspond to binaries which are already well into the LISA band when the mission starts; for large N, there is little baseline for the various modulations to encode their position.) To get a sense of more typical behavior, we selected the binaries in Figures 3-2c – 3-2i randomly from the distribution. There does not appear to be any "typical" evolution; each binary exhibits some unique features. Most binaries, however, seem to share with the binary in Figure 3-2a the property that the final day before merger gives much more information on the position than any day before it (albeit to a lesser degree). We will see below that this feature holds for the medians of almost all mass and redshift cases. It is worth noting that although KHMF agree with us on most of the other qualitative features of the time dependence, they do not see the dramatic change in the final day of inspiral. As we will discuss in more detail below, this dramatic improvement toward the end of inspiral is due to spin precession physics. This is in good agreement with the expectations of Neil Cornish [56] and KHMF that spin precession would most dramatically impact the last week or so of inspiral.

While the evolution of parameter errors for individual binaries is interesting, of more relevance is the evolution of the errors' *distribution*. The top panel of Figure 3-3 shows the time dependence of the distribution of the major axis $2a$ for our model system of $m_1 = 10^6 M_\odot$, $m_2 = 3 \times 10^5 M_\odot$, and $z = 1$. As before, data were produced for $N = 0$, 1, 2, 4, 7, 14, 21, and 28, but for clarity we only show a subset of these in the figure. We can clearly see the evolution to smaller major axis as the binary nears merger. Four weeks before merger, the median major axis is 4.8 times larger than at merger; this number shrinks to 3.9 two weeks before merger, 3.2 one week before, and 2.5 two days before. As expected from the individual binaries, the most dramatic change in the distribution occurs during the last day before merger. Not only is the median substantially reduced (by a factor of 2.2), but the shape sharply changes. For $N > 0$, the distribution is distinctly peaked, becoming gradually flatter as N gets smaller. Over the last day of inspiral, the distribution evolves into the almost entirely flat, slightly bimodal shape first seen in Figure 3-1. We find that this same behavior holds for all masses and redshift cases of interest: A sharply peaked

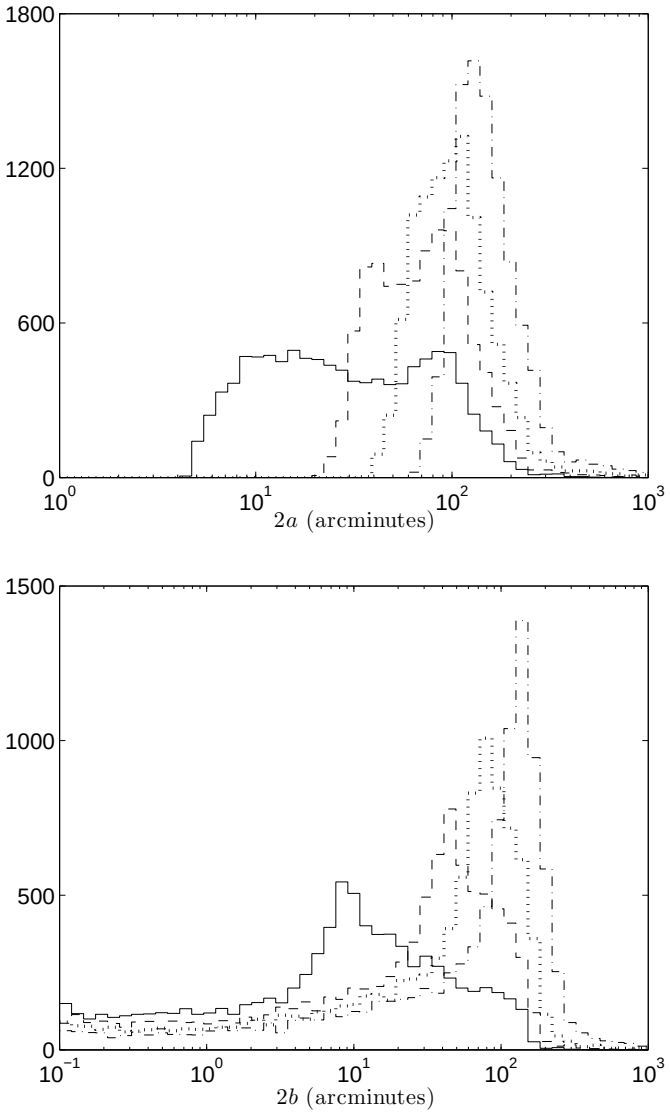

Figure 3-3: Same as Fig. 3-1, but with $m_1 = 10^6 M_\odot$, $m_2 = 3 \times 10^5 M_\odot$, and $z = 1$ at different values of N (the number of days before merger). Reading from left to right, $N = 0$ (solid line), 1 (dashed line), 7 (dotted line), and 28 (dash-dotted line). Clearly, the largest change — in shape as well as median — happens between merger and one day before.

distribution at $N > 0$ evolves into the flatter, sometimes bimodal final distributions described in Sec. 3.2. As the total mass increases, however, the final distributions become so narrow that the shape change is no longer very clear. As we might expect, this transition occurs at smaller total mass for higher z.

The bottom panel of Figure 3-3 shows the evolution of the minor axis $2b$. Again the distribution slowly changes shape over time, with the most drastic change occurring in the last day. Here the final distribution still retains a slight peak, along with the previously discussed long tail of small errors. Interestingly, this tail is present to some degree throughout the evolution. As the total mass increases, the final distribution moves to the right until, as with $2a$, the sharp change of shape disappears. The same evolution occurs at higher z, again with a shift in the mass scale.

To further understand how the error ellipse evolves, consider Figure 3-4. Here we show the evolution of the median axes for a wide range of masses ("low," "intermediate," and "high") at $z = 1$. Again, our calculations only produce output for $N = 0$, 1, 2, 4, 7, 14, 21, and 28; these points are connected by lines to guide the eye. Almost all the cases we present show similar behavior: The ellipses gradually shrink with time before sharply decreasing in size during the final day. Significant deviations from this behavior come from the high-mass binaries, which evolve more drastically at large N before settling in to resemble the lower-mass curves. This high-mass deviation is an artifact of our choice of maximum N. Smaller binaries may spend many months or even years in the LISA band before merger. In those cases, enough signal has already been measured at $N = 28$ to locate the binary reasonably well. By contrast, the high-mass binaries spend much less time in band[3] and have not been measured so well by $N = 28$. They have to "catch up" to the smaller-mass binaries over the first few weeks of our measurement window. Nearly identical results were found by KHMF: Figure 2 of KHMF plots the evolution of sky position for an intermediate-mass binary from $N \simeq 300$. They find that the measurement accuracy rapidly evolves early in the measurement, with slopes of angular error versus time very similar to what we show in Figure 3-4 for high-mass binaries.

[3]The two highest-mass binaries in Fig. 3-4 are not even in band a full 28 days.

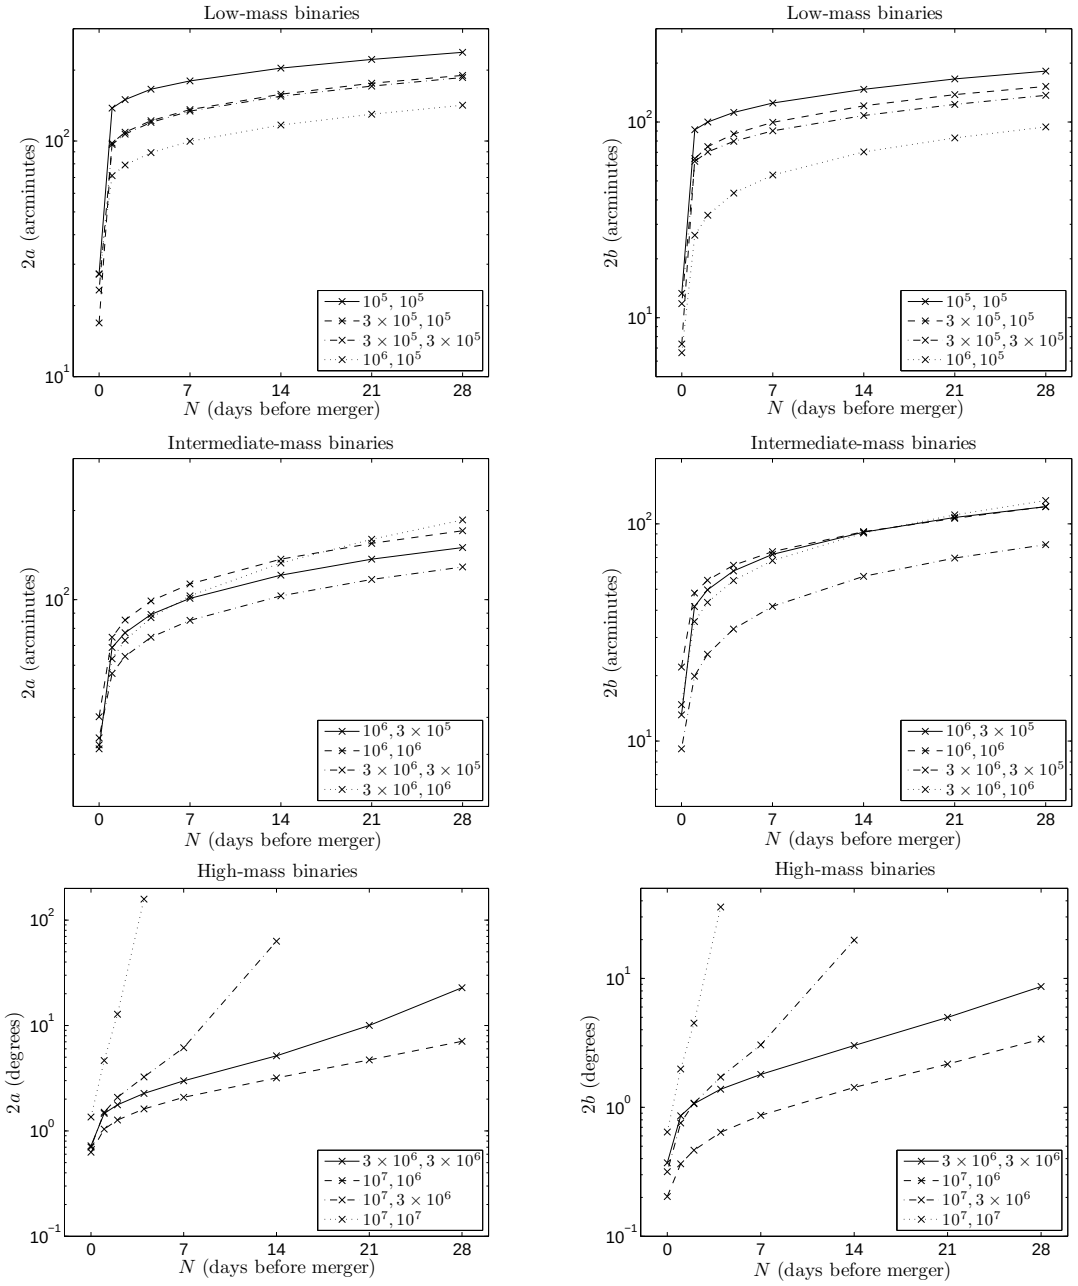

Figure 3-4: Medians of the sky position ellipse axes for Monte Carlo runs of 10^4 binaries as a function of time before merger. Major axes $2a$ are on the left; minor axes $2b$ are on the right. Data were only output at the marked points; the lines are there just to guide the eye. The masses have been subdivided into "low," "intermediate," and "high" groups; the exact values (in units of solar masses) are given in the legends. Note also the different scales (arcminutes and degrees) on the y-axis.

Although these curves are qualitatively quite similar, there are significant quantitative differences. For example, the evolution of the intermediate-mass binaries is less than that of the low-mass binaries, especially in the last day. The intermediate-mass major axes shrink by a factor of $\sim 4-6$ over the entire four-week period, whereas the low-mass axes shrink by a factor of $\sim 5-11$. Another important quantitative difference can be seen by comparing the major and minor axes. We find that the ratio $2a/2b$ grows with time in most cases, indicating that the minor axis tends to shrink more rapidly than the major axis. The only exceptions are the previously described high-mass cases, in which $2a/2b$ shrinks during most or all of the inspiral. Presumably, the same behavior would also be seen for lower-mass binaries at higher values of N; this conclusion is supported by Figure 4 of [143]. For all other cases, $2a/2b \sim 1.3-1.7$ at $N = 28$ and increases to $\sim 1.8-4$ at $N = 0$. The largest increase is typically in the final day.

What causes this dramatic improvement in the last day of inspiral? Three factors primarily contribute to our ability to localize a source on the sky: modulations due to LISA's orbital motion, modulations due to spin precession, and SNR accumulated over time. Since LISA moves the same amount in the final day as in any other day, orbital-induced modulations cannot be the cause. A great deal of SNR is accumulated in the last day (typically increasing by a factor of ~ 2 or more), and many parameter errors scale as $(\mathrm{SNR})^{-1}$. However, KHMF demonstrate that sky position and distance errors do not scale as $(\mathrm{SNR})^{-1}$ in the last few weeks before merger. Our "no precession" code supports their conclusion: The final jump in SNR cannot make up for the lack of orbital modulation over such a short time scale.

The remaining possibility is spin-induced precession. Indeed, we found in Chapter 2 that the number of modulations due to spin precession increases dramatically as the binary approaches merger (see Figure 2-1). This suggests that the improvement we see is due to the impact of precession. To examine this hypothesis, we plot the "precession" and "no precession" results together on the same axes. Figure 3-5 shows such a plot for a low-mass system and an intermediate-mass system at $z = 1$. We see that at $N = 28$ days, the two codes give similar results for localization accuracy. Their

predictions gradually diverge as merger is approached. The greatest jump between the two codes occurs on the last day before merger, agreeing with our expectation that precession effects are maximal then. The effect is greater in the low-mass case than in the intermediate-mass case. Similarly, the ratio $2a/2b$ starts about the same in both codes, growing very slowly until $N = 1$. At this point, it jumps dramatically in the "precession" code, while staying roughly the same in the "no precession" code. Interestingly, one effect of precession is that the localization errors track $(\mathrm{SNR})^{-1}$ rather closely. By breaking the various correlations which made them deviate from $(\mathrm{SNR})^{-1}$, precession-induced modulations allow the errors to evolve in a manner that is more consistent with our naive expectations.

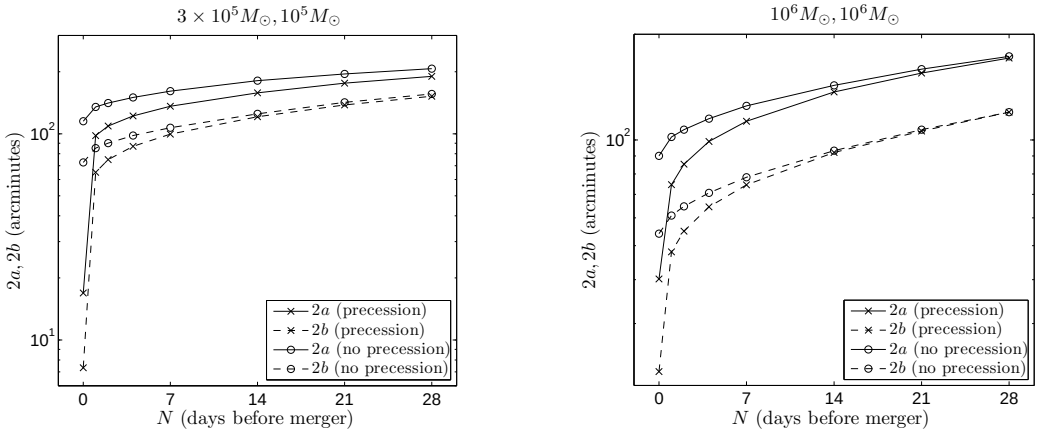

Figure 3-5: Medians of $2a$ and $2b$ as a function of time, comparing an analysis that accounts for spin-induced precession to one that neglects it. Solid lines trace the evolution of $2a$; dashed lines trace $2b$. Precession results are marked with crosses, no precession with circles. The left plot shows a low-mass case, $m_1 = 3 \times 10^5 M_\odot$ and $m_2 = 10^5 M_\odot$; the right plot shows an intermediate-mass case, $m_1 = m_2 = 10^6 M_\odot$. Both plots are for $z = 1$.

The time evolution of the correlations of interest, those between the sky position and binary orientation, are illustrated in Figure 3-6. Here we show the off-diagonal components of the covariance matrix Σ^{ab} (where $a \in \{\bar{\mu}_N, \bar{\phi}_N\}$ and $b \in \{\bar{\mu}_L(0), \bar{\phi}_L(0)\}$) for the binary in Figure 3-2a. We found that examining the normalized correlation coefficients $c^{ab} = \Sigma^{ab}(\Sigma^{aa}\Sigma^{bb})^{-1/2}$ can mislead since Σ^{aa} and Σ^{bb} are rapidly evolving at the same time as Σ^{ab}. At large N, the two sets of angles are relatively strongly correlated due to degeneracies in the measured waveform (2.45).

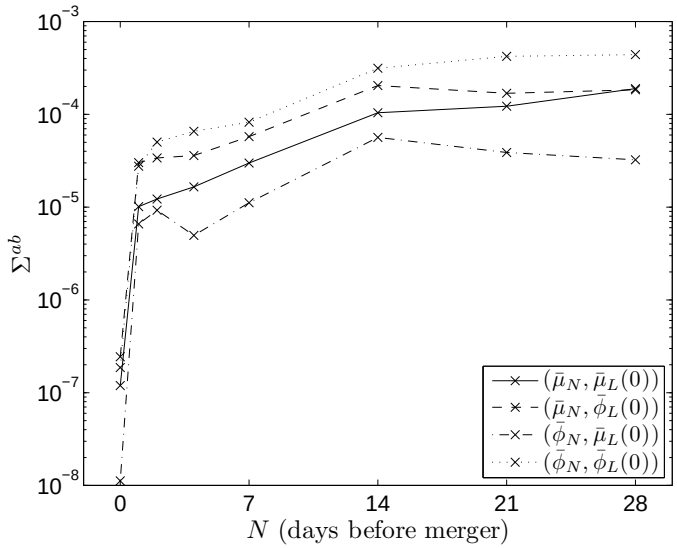

Figure 3-6: Off-diagonal covariance matrix entries illustrating correlation between sky position and binary orientation as a function of time, for the binary in Fig. 3-2a. The correlations decrease rapidly in the final day before merger, when precession effects are maximal.

However, in the last day before merger, the correlations sharply decrease as precession effects accumulate. The reduction of these correlations coincides with the sudden drop in parameter errors seen in Figure 3-2*a* (and, by extension to the entire Monte Carlo run, Figs. 3-3 – 3-5).

Finally, we investigate the time evolution of errors in the luminosity distance D_L. Figure 3-7 shows the distribution of $\Delta D_L/D_L$ (determined solely by taking into account GW measurement effects) evolving in time for a binary with $m_1 = 10^6 M_\odot$, $m_2 = 3 \times 10^5 M_\odot$, and $z = 1$. We see that in contrast to sky position, the shape of the distribution does not change very much with time. It typically spreads enough to reduce its height, but it maintains a well-defined peak. However, the progression of the median is very similar to the sky position case: It decreases slowly with time until the last day, when it jumps drastically. The evolution of median values of $\Delta D_L/D_L$ follows tracks very similar in shape to those shown in Figure 3-4, so we do not show them explicitly.

Tables 3.4–3.6 summarize all of our results on the time evolution of localization.

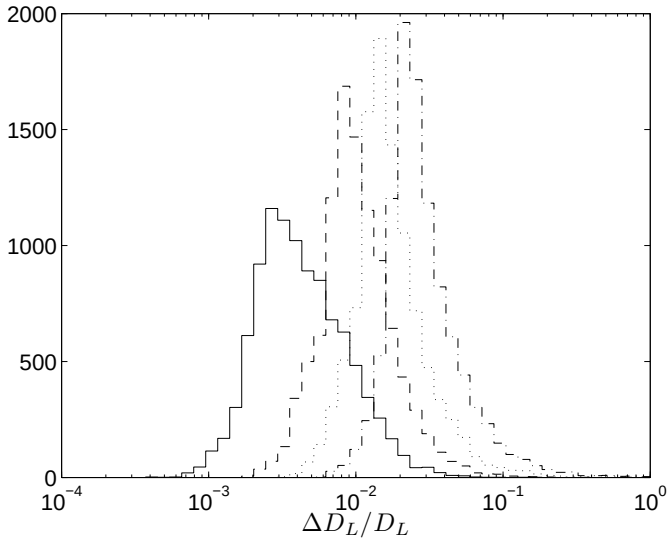

Figure 3-7: Distribution of $\Delta D_L/D_L$ for 10^4 binaries with $m_1 = 10^6 M_\odot$, $m_2 = 3 \times 10^5 M_\odot$, and $z = 1$ at different values of N (the number of days before merger). Reading from left to right, $N = 0$ (solid line), 1 (dashed line), 7 (dotted line), and 28 (dash-dotted line).

At low redshift, the ability to locate an event on the sky is quite good over much of the mass range even as much as a month in advance of the final merger. In most cases, the localization ellipse at $z = 1$ is never larger than about 10 deg^2 in size, which is comparable to the field of view of proposed future surveys, such as LSST [246]. This ability degrades fairly rapidly as redshift increases, especially for larger masses. At $z = 3$, the ellipse can be ~ 10 deg^2 a few days in advance of merger for small and intermediate masses. At the highest masses, GWs provide very little localization information. Going to $z = 5$ makes this even worse; an ellipse of ~ 10 deg^2 can be found at most a day prior to merger, and only for relatively small mass ranges.

In many cases, $\Delta D_L/D_L$ is determined so well by GWs that gravitational lensing errors are expected to dominate. As such, the GW-determined values of $\Delta D_L/D_L$ are essentially irrelevant for locating these binaries in redshift space; lensing will instead determine how well redshifts can be measured. However, in some cases, the intrinsic distance error exceeds the lensing error, so it is worth knowing when the "lensing limit" can be achieved. At $z = 1$, the limit is achieved for most binaries as long as

Mass range	N (days)	$2a$ (deg)	$2b$ (deg)	$\Delta\Omega_N$ (deg^2)	$\Delta D_L/D_L$
Low	0	$(17-27)^{\rm a}$	$(6.6-13)^{\rm a}$	$0.02-0.07$	$(2.4-4)\times10^{-3}$
	1	$(71-140)^{\rm a}$	$(26-91)^{\rm a}$	$0.3-2.8$	$(8-21)\times10^{-3}$
	7	$(100-180)^{\rm a}$	$(54-120)^{\rm a}$	$0.9-5.1$	$(1.3-3)\times10^{-2}$
	28	$(140-240)^{\rm a}$	$(94-180)^{\rm a}$	$2.6-9.6$	$(2-4.3)\times10^{-2}$
Intermediate	0	$(31-40)^{\rm a}$	$(9.2-22)^{\rm a}$	$0.04-0.18$	$(3.8-5.6)\times10^{-3}$
	1	$(56-75)^{\rm a}$	$(20-48)^{\rm a}$	$0.17-0.77$	$(6.4-11)\times10^{-3}$
	7	$(85-110)^{\rm a}$	$(42-74)^{\rm a}$	$0.6-1.9$	$(1.1-1.7)\times10^{-2}$
	28	$(130-190)^{\rm a}$	$(80-130)^{\rm a}$	$2-5.4$	$(1.8-3)\times10^{-2}$
High	0	$0.6-1.4$	$0.2-0.64$	$0.07-0.68$	$(4.6-25)\times10^{-3}$
	1	$1-4.7$	$0.4-2$	$0.2-7.1$	$(0.7-16)\times10^{-2}$
	7[b]	$2.1-6.2$	$0.9-3$	$1-14$	$(1.6-8.7)\times10^{-2}$
	28[b]	$7.1-23$	$3.4-8.7$	$16-150$	$(8.4-30)\times10^{-2}$

Table 3.4: Typical ranges of sky position and distance measurement accuracy as a function of time until merger for binaries at $z = 1$. The total mass M has been divided into three categories: low mass ($M \lesssim 10^6 M_\odot$), intermediate mass ($10^6 M_\odot \lesssim M \lesssim 4\times10^6 M_\odot$), and high mass ($M \gtrsim 6\times10^6 M_\odot$). Angles are in degrees except for those marked "a," which are in arcminutes; solid angles are always in square degrees. For rows marked "b," some very massive systems are excluded from the data. In those cases, the position and distance are very poorly constrained that far in advance of merger. In some cases, the binary is even out of band.

a month before merger. For $z = 3$, the limit is only achieved a few days to a week (depending on mass) before merger; for $z = 5$, intrinsic errors generally exceed the lensing errors even at a day before merger.

3.3.3 Results II: Angular dependence of localization accuracy

We now examine one final interesting property of the errors: their dependence on the sky position of the source. As we design future surveys to find counterparts to MBHB coalescences, it will be important to understand if there is a bias for good (or bad) localization in certain regions of the sky. It is also useful to know in advance whether the "best" regions are likely to be blocked by foreground features such as

Mass range	N (days)	$2a$ (deg)	$2b$ (deg)	$\Delta\Omega_N$ (deg^2)	$\Delta D_L/D_L$
Low	0	$1.3-2.4$	$0.6-1.3$	$0.7-2.1$	$(1.2-2)\times10^{-2}$
	1	$3.9-7.2$	$1.4-4.7$	$3-27$	$(2.7-6.3)\times10^{-2}$
	7	$5.8-9.6$	$2.9-6.5$	$10-51$	$(4.4-9.3)\times10^{-2}$
	28	$8.5-13$	$5.4-9.8$	$32-100$	$(7.2-14)\times10^{-2}$
Intermediate	0	$2.1-2.6$	$0.7-1.4$	$0.8-2.6$	$(1.5-2.4)\times10^{-2}$
	1	$3.6-5$	$1.3-3.1$	$2.4-12$	$(2.5-4.5)\times10^{-2}$
	7	$6-11$	$2.8-6.3$	$9.8-48$	$(4.5-10)\times10^{-2}$
	28	$9.8-91$	$6.3-42$	$43-2900$	$(0.9-14)\times10^{-1}$
High	0	$3.4-33$	$1.3-9.7$	$2.7-260$	$(2.9-240)\times10^{-2}$
	1[a]	$5.7-17$	$2.4-7.3$	$8.5-93$	$(5.2-53)\times10^{-2}$
	7[a]	$25-75$	$13-27$	$220-1500$	$(5.8-19)\times10^{-1}$
	28[b]	—	—	—	—

Table 3.5: Typical ranges of sky position and distance measurement accuracy as a function of time until merger for binaries at $z = 3$. The total mass M has been divided into three categories just as in Table 3.4. For rows marked "a," some very massive systems are excluded from the data. In those cases, the position and distance are very poorly constrained that far in advance of merger. In some cases, the binary is even out of band. For the row marked "b," all of the binaries of that mass range are either very poorly measured or completely out of band so far in advance of merger.

the Galactic center.

Before discussing this dependence in detail, it is worth reviewing some details of how our Monte Carlo distributions are constructed. As described in Chapter 2, in most of our analysis we randomly distribute the sky position of our binaries, drawing from a uniform distribution in $\bar{\mu}_N = \cos\bar{\theta}_N$ and $\bar{\phi}_N$, where $\bar{\theta}_N$ and $\bar{\phi}_N$ are the polar and azimuthal angles of a binary in solar system barycenter coordinates. We also randomly choose our binaries' final merger time. In this section, rather than distributing the sky position, we examine parameter accuracies for particular given positions; all other Monte Carlo parameters are distributed as usual. Because we continue to randomly distribute the final merger time, however, the *relative* azimuth between a binary's sky position and LISA's orbital position at merger, $\Delta\phi = \bar{\phi}_N - \bar{\Phi}_D(t_c)$, remains randomly distributed. As such, we expect our analysis to effectively average over $\bar{\phi}_N$, washing out any strong dependence on this angle.

Mass range	N (days)	$2a$ (deg)	$2b$ (deg)	$\Delta\Omega_N$ (deg^2)	$\Delta D_L/D_L$
Low	0	$2.8-4.9$	$1.1-2.7$	$2.3-9.3$	$(2.6-4.1)\times10^{-2}$
	1	$6.8-12$	$2.4-7.6$	$9-71$	$(4.6-10)\times10^{-2}$
	7	$10-16$	$5-11$	$32-140$	$(7.8-15)\times10^{-2}$
	28	$16-22$	$9.6-16$	$100-290$	$(1.3-2.3)\times10^{-1}$
Intermediate	0	$3.9-5.2$	$1.4-2.7$	$3.3-10$	$(3.2-5)\times10^{-2}$
	1	$7.2-11$	$2.5-6.2$	$9.6-48$	$(5.2-10)\times10^{-2}$
	7	$12-41$	$5.9-21$	$45-610$	$(1.1-5.2)\times10^{-1}$
	28[a]	$21-170$	$15-65$	$250-8000$	$(2-21)\times10^{-1}$
High	0[a]	$9-29$	$4.3-10$	$30-230$	$(1.2-12)\times10^{-1}$
	1[a]	$19-27$	$9.9-12$	$130-250$	$(4.4-7.3)\times10^{-1}$
	7[b]	—	—	—	—

Table 3.6: Typical ranges of sky position and distance measurement accuracy as a function of time until merger for binaries at $z = 5$. The total mass M has been divided into three categories just as in Table 3.4. For rows marked "a," some very massive systems are excluded from the data. In those cases, the position and distance are very poorly constrained that far in advance of merger. In some cases, the binary is even out of band. For the row marked "b," all of the binaries of that mass range are either very poorly measured or completely out of band so far in advance of merger.

We begin by examining the dependence of errors on $\bar{\mu}_N$. We evenly divide the range $-1 \leq \bar{\mu}_N \leq 1$ into 40 bins and run a Monte Carlo simulation with 10^4 points in each. That is, we pick $\bar{\mu}_N$ only from the bin range, but we pick all other random parameters in the usual manner. The results for a representative binary ($m_1 = 10^6 M_\odot, m_2 = 3\times10^5 M_\odot$, and $z = 1$) are shown in Figure 3-8.

Note that all of the error distributions are symmetrically peaked around the plane of LISA's orbit ($\bar{\mu}_N = 0$). Any slight asymmetry is due only to statistical effects. This is reassuring; LISA should not favor one hemisphere over the other. There is also additional structure that is parameter dependent. At its peak, the major axis $2a$ is almost 35% greater than the position-averaged median value[4] of $31.1'$. It then decreases with $|\bar{\mu}_N|$ and reaches a minimum of about $25'$ for $0.75 < |\bar{\mu}_N| < 0.8$. Finally, there are subpeaks near the ecliptic poles, although they still lie below the

[4]Note that the position-averaged medians quoted here are slightly different from those quoted in Tables 2.4–2.9, since this sample has 40 times more points.

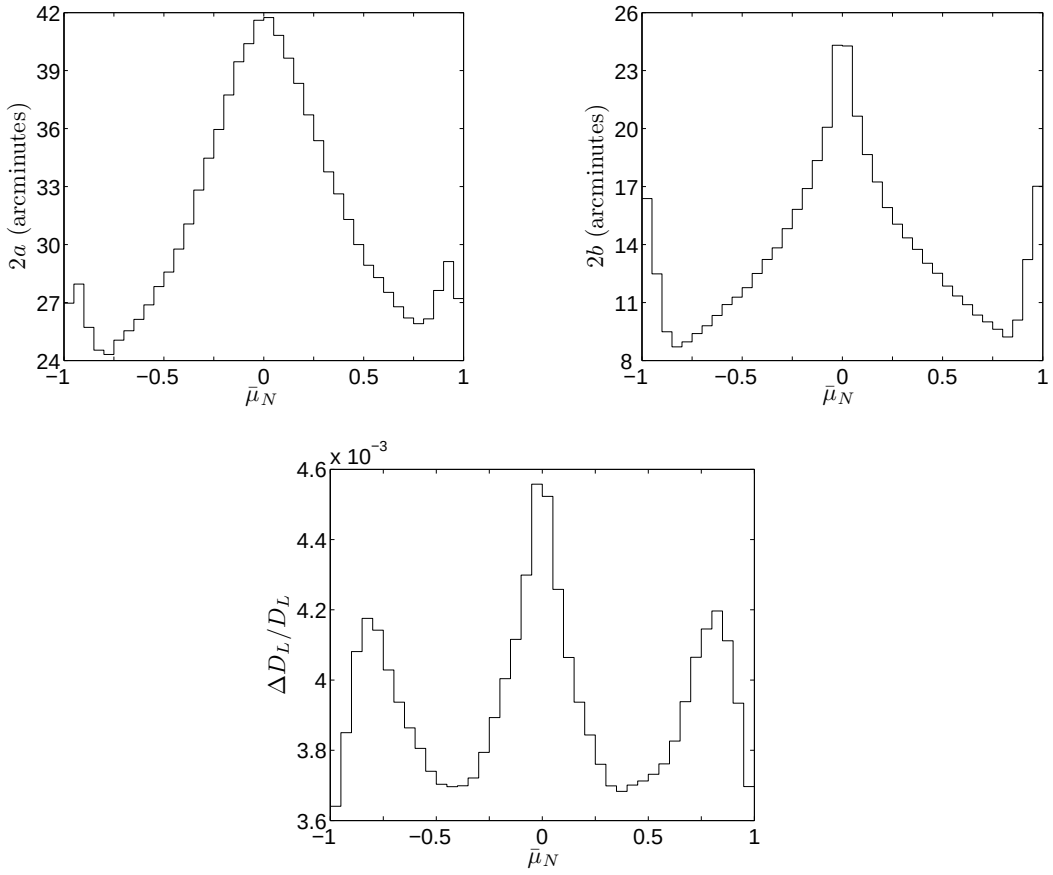

Figure 3-8: Dependence of the localization errors on $\bar{\mu}_N$. The major axis $2a$ of the sky position error ellipse is on the top left, the minor axis $2b$ on the top right, and the luminosity distance errors $\Delta D_L/D_L$ on the bottom. Each datum represents the median of 10^4 binaries with $m_1 = 10^6 M_\odot$, $m_2 = 3 \times 10^5 M_\odot$, and $z = 1$; all other parameters are selected randomly (except for $\bar{\mu}_N$, whose range is limited to the bin width).

position-averaged median. The dependence of the minor axis $2b$ on angle is even more dramatic. At its peak, $2b$ differs from the position-averaged median of $13.0'$ by over 85%. Just like the major axis, it drops to a minimum, but it does so more rapidly. The minimum also occurs at a slightly larger value of $|\bar{\mu}_N|$ than for $2a$. The minor axis also shows fairly strong subpeaks near the ecliptic poles, with values higher than the position-averaged median.

The luminosity distance errors behave slightly differently. First of all, the variation with $\bar{\mu}_N$ is weaker than for the sky position: At the central peak, $\Delta D_L/D_L$ is only $\sim 15\%$ greater than its position-averaged median value (0.00392). In addition, while the distribution again peaks near the poles, it does so at a smaller value of $|\bar{\mu}_N|$. (In fact, the peaks occur very close to where sky position errors are *minimized*.)

By binning the data sets developed for Chapter 2, we are able to confirm this behavior over a wide range of masses and redshifts, albeit with poorer statistics: 10^4 binaries in total, rather than per bin. Thanks to the poorer statistics, we cannot resolve the small polar subpeaks in the distribution of $2a$. In addition, in some cases it appears that the side peaks can be larger than the central peak in the distribution of $\Delta D_L/D_L$. Aside from these minor variations, the shapes and relative amplitudes seen in these distributions hold robustly over all masses and redshifts we consider.

To examine the role that precession plays in determining the dependence of localization accuracy on $\bar{\mu}_N$, we perform the same analysis using the "no precession" code. We find similar distributions for $2a$ and $2b$ (though somewhat broader); we also find that $2a$ more closely resembles $2b$ at the poles. $\Delta D_L/D_L$ is quite a bit different: It is *minimized* in the ecliptic plane and increases monotonically to the poles. Finally, careful checking shows that the distributions including precession physics at 28 days before merger have the same shape as the distributions neglecting precession do at merger; they then evolve to those in Fig. 3-8 over time.

We next investigate the dependence of the errors on the azimuthal angle $\bar{\phi}_N$. To improve the statistics, we now calculate 10^5 binaries in each bin; the results are shown in Figure 3-9. The errors have a very weak (although nonzero) dependence on $\bar{\phi}_N$: The maximum deviation from the overall median is only $\sim 1\% - 2\%$. This is to

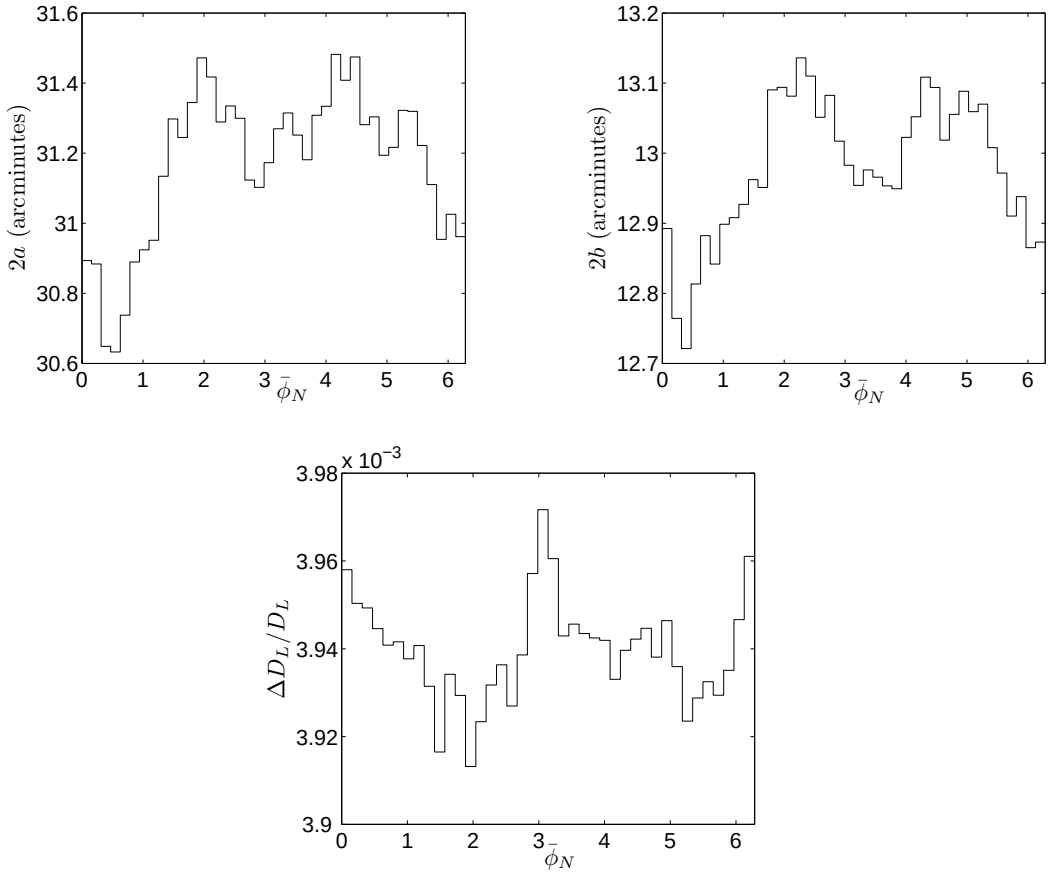

Figure 3-9: Dependence of the localization errors on $\bar{\phi}_N$. The major axis $2a$ of the sky position error ellipse is on the top left, the minor axis $2b$ on the top right, and the luminosity distance errors $\Delta D_L/D_L$ on the bottom. Each datum represents the median of 10^5 binaries with $m_1 = 10^6 M_\odot$, $m_2 = 3 \times 10^5 M_\odot$, and $z = 1$; all other parameters are selected randomly (except for $\bar{\phi}_N$, whose range is limited to the bin width).

be expected; as discussed at the beginning of this section, the randomness of our binaries' merger times effectively averages over azimuth. If we did not average over azimuth in this way, we would expect a moderately strong $\bar{\phi}_N$-dependence due to the functional form of LISA's response to GWs. Even after averaging this dependence away, we might expect some residual $\bar{\phi}_N$ structure due to the "rolling" motion of the LISA constellation. The phase associated with LISA's roll angle puts an additional oscillation on the measured waves (see the α-dependence in Eqs. (47) and (48) of KHMF, where α encodes the roll angle), and we do not average over this angle. Indeed, on close inspection, we can make out roughly two peaks in each plot in Figure 3-9 (although the statistics are still too poor to resolve them clearly), consistent with the $\cos 2\alpha$ and $\sin 2\alpha$ behavior shown in KHMF. For comparison, we also examine the $\bar{\phi}_N$-dependence for this mass and redshift with precession turned off. The oscillatory behavior appears very clearly in this case. Because of the weakness of the $\bar{\phi}_N$ behavior, we are unable to easily check it for other masses and redshifts.

Finally, to give an overall sense as to how localization varies on the sky, we present in Figure 3-10 sky maps of the median major axis $2a$, minor axis $2b$, localization ellipse area $\Delta\Omega_N$, and distance accuracy $\Delta D_L/D_L$. We show these data both in ecliptic and Galactic coordinates. Note that most of the region of small error lies outside of the Galactic plane. This potentially bodes well for searches for MBHB electromagnetic counterparts — the regions where instruments like LISA "see" most sharply are less likely to be hidden by foreground features. Certain portions of the sky that LISA sees well will be easier to search telescopically than others. It will be an important task for future surveys over all electromagnetic bands to identify regions that are particularly amenable to finding counterparts to MBHB events.

3.4 Summary and conclusions

As discussed at length in Chapter 2, accounting for the general relativistic precession of the angular momentum vectors in an MBHB system has a dramatic impact on what we can learn by observing the system's gravitational waves. Spin-induced precession

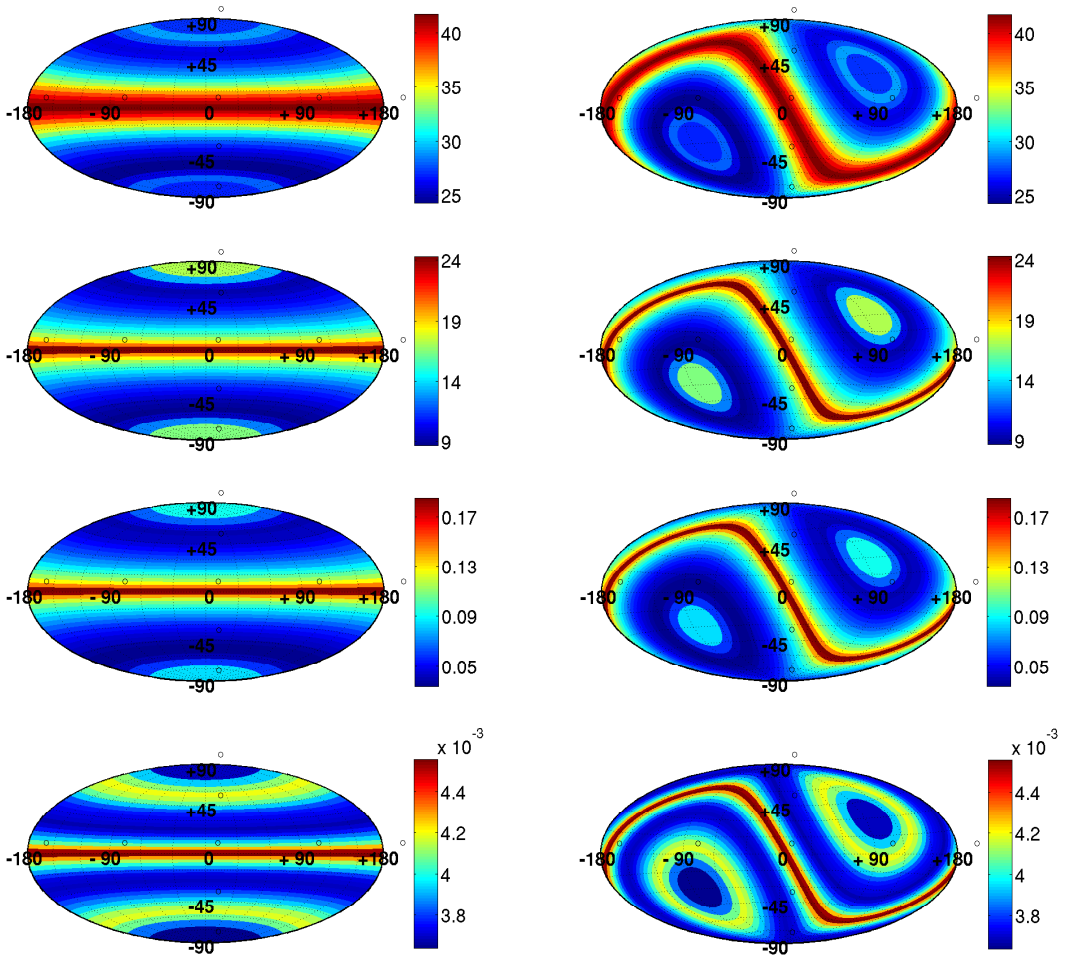

Figure 3-10: Sky maps of major axis $2a$ (top row, in arcminutes), minor axis $2b$ (second row, in arcminutes), localization ellipse area $\Delta\Omega_N$ (third row, in square degrees) and $\Delta D_L/D_L$ (bottom row) for LISA observations of MBH binaries in different parts of the sky. Data in the left column are presented in ecliptic coordinates; data on the right are in Galactic coordinates, with the Galactic center at the middle. Note that the level of $\bar{\phi}_N$ variation is so small that it would not show up in these figures; accordingly, they are essentially just remappings of Fig. 3-8.

breaks degeneracies among different parameters, making it possible to measure them more accurately than they could be determined if precession were not present. This has a particularly important impact on our ability to locate such a binary on the sky and to determine its luminosity distance — the degeneracy between sky angles, distance, and orientation angles is severe in the absence of precession.

Our analysis shows that the improvement that precession imparts to measurement accumulates fairly slowly. In using one code which includes the impact of spin precession and a second which neglects this effect, we find little difference in the accuracy with which GWs determine sky position and distance for times more than a few days in advance of the final merger. The difference between the two codes grows quite rapidly in these final days. In the last day alone, the localization ellipse area decreases by a factor of $\sim 3-10$ (up to ~ 60 in a few low-mass systems) when precession effects are included. Distance determination is likewise improved by factors of $\sim 1.5-7$ in that final day.

Not all of the precession effects occur in the final days. Close examination of Figure 3-5 shows that for larger mass ratios, at least some improvement from precession can be seen in advance of merger. More strikingly, we saw in Figure 3-3 that the long tail of small minor axes can be seen, to some degree, throughout the inspiral. We could get lucky and find a binary with a very small value of $2b$ weeks before merger. But the large improvement in the median that we found in Chapter 2 appears to take effect only in the final days of inspiral. Therefore, while precession may in fact help improve the *final* localization of a coalescing binary by a factor of $\sim 2-10$ in each direction, it will not be too much help in *advanced* localization of a typical binary.

Nevertheless, the pixel sizes that we find are small enough that future surveys should not have too much trouble searching the region identified by GWs, at least over certain ranges of mass and redshift. At $z = 1$, the GW localization ellipse is ~ 10 deg^2 or smaller for most binaries as early as a month in advance of merger. (At high masses, the ellipse can be substantially larger than this a month before merger, but it shrinks rapidly, reaching a comparable size $1-2$ weeks before merger.) This bodes well for future surveys with large fields of view that are likely to search the GW

pixel for counterparts. In addition, GWs determine the source luminosity distance so well that the distance errors we find are essentially irrelevant — gravitational lensing will dominate the distance error budget for all but the highest masses.

As redshift increases, the GW pixel rapidly degrades, particularly for the largest masses. Let us adopt 10 deg^2 (the approximate LSST field of view) as a benchmark localization for which counterpart searches may be contemplated. At $z = 3$, this benchmark is reached at merger for almost the entire range of masses we considered. As little as a day in advance of merger, however, some of the least massive and most massive systems are out of this regime. One week prior to merger, the most massive systems are barely located at all (ellipses hundreds of square degrees or larger). The intermediate masses do best, but even in their cases the positions are determined with ~ 10 deg^2 accuracy no earlier than a few days in advance of merger. The resolution degrades further at higher redshift. At $z = 5$, systems with $M \gtrsim 6 \times 10^6\, M_\odot$ are not located more accurately than ~ 30 deg^2 even at merger. Smaller systems are located within ~ 10 deg^2 at merger, but very few are at this accuracy even one day in advance of merger. The luminosity distance errors also increase, so much that they exceed lensing errors a few days to a week before merger at $z = 3$, and only a day before merger at $z = 5$. This degradation hurts the ability to search for counterparts by redshift and subsequently use them as standard candles.

Our main conclusion is that future surveys are likely to have good advanced knowledge (a few days to one month) of the location of MBHB coalescences at low redshift ($z \sim 1 - 3$), but only a day's notice at most at higher redshift ($z \sim 5$). This conclusion may be excessively pessimistic. As mentioned earlier, recent work examining the importance of subleading harmonics of MBHB GWs is finding that including harmonics beyond the leading quadrupole has an important effect on the final accuracy of position determination [16, 245, 196]. For most masses, these analyses show a factor of a few improvement in position, comparable to the improvement that we find when spin precession is added to the waveform model. For high-mass systems, the higher harmonics increase the (previously small) overlap with the LISA band; consequently, the improvement can be much larger, up to 2 or 3 orders of magnitude

in area. Even though preliminary results indicate that these improvements will not combine with those from precession, the case is far from complete. In particular, the time dependence of the higher harmonic improvement is unknown.

Finally, we have also studied the sky position dependence of LISA's ability to localize sources. We have found that the regions of best localization lie fairly far out of the Galactic plane. However, as emphasized by Neil Cornish [57], a proper anisotropic confusion background might impact this dependence. In our calculations, we have assumed an isotropic background, neglecting the likely spatial distribution of Galactic binaries. Properly accounting for this background is likely to strengthen our conclusion that LISA's ability to "see" is best for MBHB sources out of the Galactic plane.

Chapter 4

Tidal perturbations of neutron stars

4.1 Introduction

4.1.1 Neutron stars and the equation of state

Neutron stars have a very complicated structure [221, 47]. The outside of a neutron star, called the *outer crust*, typically consists of matter at relatively "low" density ($10^6\,\mathrm{g\ cm^{-3}} < \rho < 4 \times 10^{11}\,\mathrm{g\ cm^{-3}}$), in the form of heavy nuclei (e.g., iron, the end point of nuclear fusion in the star's progenitor) and relativistic, degenerate electrons. As distance from the center decreases and density increases, the nuclei become more and more neutron-rich as electrons and protons combine to form neutrons (inverse beta decay). At $\rho \sim 4 \times 10^{11}\mathrm{g\ cm^{-3}}$, the density of "neutron drip," free neutrons appear. In the *inner crust* ($4 \times 10^{11}\,\mathrm{g\ cm^{-3}} < \rho < 2 \times 10^{14}\,\mathrm{g\ cm^{-3}}$), therefore, a superfluid neutron gas exists along with the heavy nuclei and electrons. The *interior* of the star begins above nuclear density, $\rho \sim 2 \times 10^{14}\,\mathrm{g\ cm^{-3}}$. Here the nuclei break apart and the star is composed primarily of superfluid neutrons. This model, a star made of free neutrons, is that proposed originally by Baade and Zwicky [17]. On the most basic level, this star can be said to be supported by the degeneracy pressure of neutrons, much like a white dwarf is supported by the degeneracy pressure of

electrons.

Unfortunately, the reality is not so simple. In particular, at such high densities, the interactions between nucleons become very important. These interactions are difficult to model, especially when the ratio of neutrons to protons is large [162]. The best way to describe the nature of matter is via the *equation of state* (EOS), a relation between pressure and density.[1] The equation of state of degenerate matter, like that in a white dwarf, can be expressed very accurately by a *polytrope*, an equation of state in which the pressure varies as some power of the density. However, the equation of state of strongly interacting neutron-rich matter, especially with superfluid properties, is much more complicated. In addition, at densities a few times nuclear density, more exotic effects might occur. Examples include the excitations of hyperons (baryons containing strange quarks), Bose condensation of mesons (pions or kaons), or dissociation into free quarks. It is even possible that the entire star could be made of strange quark matter.

The uncertainty about the core equation of state means that neutron stars can be used as *laboratories* to study the nature of matter in extremely dense, low-temperature conditions, a regime not well probed by experiments on Earth (e.g., colliders). In particular, the equation of state affects bulk properties of the star, such as the mass $m_{\rm NS}$ and the radius R. For each potential equation of state, the relativistic stellar structure equations can be solved to produce a curve in the mass-radius plane. Then actual measurements of the mass and radius can be used to place constraints and ultimately identify the correct equation of state. (Similar plots can be made for other observables, such as the mass and moment of inertia.)

Neutron star masses are easily measured by binary dynamics. (See [163] for a relatively current summary of known masses.) The most accurate measurements come from binary pulsars, in which relativistic effects help resolve the inclination angle degeneracy. These masses are generally clustered very strongly around 1.4 $M_\odot$. Masses measured in neutron star-white dwarf binaries or X-ray binaries can generally

[1] Technically, this is a *barotropic* equation of state, in which the pressure does not depend on the temperature. This is a good approximation for neutron stars, which have $T \ll T_{\rm Fermi}$.

be much larger, up to and exceeding $\sim 2\ M_\odot$; however, in the largest mass cases, the error bars are too large to make any convincing statements. The smallest well-measured mass is $1.18 \pm 0.02\ M_\odot$ [85]. Other stars reach $1\ M_\odot$ or below, but with large error bars [163].

Neutron star radii are much more difficult to determine. A simple constraint [163] is that $R > 2m_{\rm NS}$; otherwise, the object would be a black hole. It turns out that for finite pressure, $R > (9/4)m_{\rm NS}$; this statement, known as Buchdahl's theorem, can be derived from the relativistic structure equations given below. The requirement of causality, that the sound speed be less than the speed of light, sets the limit $R > 2.9m_{\rm NS}$. On the other hand, a maximum radius can be obtained by considering the break-up limit for rotating stars; the minimum period is $P_{\rm min} \sim \sqrt{R^3/m_{\rm NS}}$. The fastest rotating neutron stars are the millisecond pulsars, pulsars which were spun up by accretion from a binary companion. Lattimer and Prakash [163] plot the rotation limit for the 716 Hz pulsar J1748-2446ad [114].

Measuring the thermal emission from a neutron star's surface gives a quantity called the radiation radius, a combination of the mass and radius [162, 163]. The mass enters the expression due to the gravitational redshift of the photons.[2] The redshift itself can be measured from spectral lines. Together, these two quantities could give $m_{\rm NS}$ and R separately. Radiation radius measurements are made difficult by the complexities of neutron star atmospheres (including magnetic effects) and uncertainty in the distance to the star. For sources with periodic thermonuclear X-ray bursts separated by quiescent thermal emission, the Eddington flux during the bursts sets another constraint, allowing $m_{\rm NS}$ and R to be determined without the distance.

Other methods also exist [162, 163]. For instance, measurements of quasi-periodic oscillations (QPOs) in low-mass X-ray binaries give the orbital frequency of accreting matter at the innermost stable circular orbit (ISCO). From this frequency, upper limits can be found on $m_{\rm NS}$ and R. In the double pulsar PSR J0737-3039A & B [166],

[2]This effect, in which photons lose energy climbing out of a gravitational field, is one of the simplest consequences of the Equivalence Principle.

spin precession effects like those discussed in Chapter 2 can be measured, setting limits on the moments of inertia. Physics of the star's crust can be probed by pulsar glitches (momentary spin-ups), "starquakes" after soft gamma-ray repeater flares, and thermal relaxation in between X-ray bursts.

Finally, the NS equation of state may be constrained by tidal effects in compact binaries. These effects can manifest in the gravitational wave signals from these binaries. One potential source of information is *tidal disruption* in neutron star-black hole systems. It has been suggested that Advanced LIGO might be able to provide information on the NS equation of state from tidal disruption waves [248]. The exact form of the waveform during disruption is likely to be quite complicated, but the disruption event should "shut off" the waves in an obvious way at a certain frequency which is a function of the equation of state. More detailed calculations of the disruption waveform in numerical relativity could lead to even more information. Even if the neutron star does not tidally disrupt, tidal *distortions* could still impact the inspiral waveform. Post-Newtonian waveforms like those used in Chapters 2 and 3 traditionally do not take these finite-size effects into account. Even though they formally enter at 5PN order (i.e., $\sim v^{10}$), the tidal terms turn out to be larger than the (unknown) point-particle 5PN terms by a factor of $(R/m_{\rm NS})^5$, which can be quite large for a neutron star [95]. For some equations of state, these terms might even be large enough to measure.

4.1.2 Outline of this chapter

This chapter presents calculations of the tidal disruption and distortion of neutron stars in compact binaries. As mentioned above, such effects could impact the gravitational wave signals from these binaries, providing information on the equation of state. It is also interesting to know when a neutron star tidally disrupts because black hole-neutron star binaries are thought to be strong candidates for short gamma-ray bursts. If the star tidally disrupts outside the innermost stable circular orbit, the material can form a stable accretion disk which may lead to the burst phenomenon. If, instead, the star tidally disrupts within the ISCO, the disk will not form.

Of course, all tidal effects are most accurately modeled using numerical relativity. Although that field has developed rapidly in the past several years (see [225, 226, 232, 80, 224, 203, 260, 81, 223] for recent studies of BH-NS systems), numerical simulations remain computationally expensive, especially those involving matter. A quasi-analytic method, while only approximate, could investigate large areas of parameter space quickly, identifying interesting situations for follow-up with full numerical calculations.

Section 4.2 of this chapter reviews the history of the tidal disruption problem. All previous studies treat the disrupting star's self-gravity as Newtonian, but this approach is clearly inappropriate for a neutron star. The goal of the first part of this chapter is to go *beyond* this limitation and include general relativistic self-gravity in the disruption model. The method we use is perturbation theory. Section 4.3 describes this method in detail. First, we present the standard Tolman-Oppenheimer-Volkoff (TOV) structure equations for an unperturbed star. We then discuss the general nonradial perturbations to a spherically symmetric spacetime and derive the specific equations governing the static perturbations to a neutron star. The perturbations are *matched* to a specific black hole tidal field in Sec. 4.4.

Section 4.5 presents a few more details of our method and then discusses the results. Our results are compared to three benchmarks: (1) a previous study which used Newtonian self-gravity [254] (on which we are trying to improve), (2) numerical relativity results [232] (the correct answer), and (3) a recent study which implements relativistic self-gravity in a different way [88]. We find that our stars typically disrupt at a radius $\sim 75\%$ smaller than the numerical results, while the results of [88] match the numerical results quite well. To investigate the validity of our method further, we use it to compute the tidal distortion and disruption of a white dwarf, for which Newtonian self-gravity is appropriate. These results are then compared to the traditional calculations of Sec. 4.2. We find that deviations from the true result can be attributed both to the use of small perturbations in a strong-field regime and the neglect of rotation and fluid motion effects.

Since our tidal disruption results turn out to be not as accurate as hoped, we

change gears in Sec. 4.6 to examine small distortions in a neutron star caused by a *general* tidal field. These distortions can be characterized by a quantity known as the Love number. It turns out that the tidal oscillations generated by the black hole have an impact on the gravitational wave phase that depends, to lowest order, on the Love number. In Sec. 4.6, we derive this effect (following [95]) and see how the Love number emerges quite easily from our results of Sec. 4.3.

Section 4.7 presents the results, beginning with the Love number for a polytrope equation of state. This was previously calculated by Hinderer [116], but we correct an error in her equations that led to an underestimate of the relativistic effects. Then we extend the work of Hinderer by calculating the relativistic Love numbers of various realistic equations of state. We also present the total number of gravitational wave cycles generated by the tidal distortion term throughout the band of Advanced LIGO. We find that for two 1.4 $M_\odot$ neutron stars in a binary, only a few equations of state change the GW phase by more than 1 cycle. The most important factor seems to be the radius R for a given equation of state; large R stars have an enhanced tidal effect. For lower masses, the effect is also greater.

Finally, we conclude in Sec. 4.8 by discussing possible avenues for future research.

4.2 Newtonian models of tidally distorted stars

We seek to model the structure of a fluid star (a neutron star) in the tidal field of another mass (a black hole). The problem was first considered by Roche in 1847-50, who investigated the equilibrium structure of a satellite orbiting a rigid body [208]. In his scheme, the distorted satellite is "tidally locked" to the main body; that is, its rotation speed is exactly equal to its orbital velocity. Roche found that at a certain orbital radius, no equilibrium solution exists. He considered the body to be *tidally disrupted* at this point, called the Roche limit. Note that this formalism turns a fundamentally *dynamic* problem into a *static* problem (existence of equilibrium solutions); we shall do the same in our relativistic formalism.

In 1906, Darwin attempted to generalize the Roche problem by removing the

restriction that the second body be rigid, allowing it to tidally deform as well [66]. Later, in 1917, Jeans studied a simpler form of the Roche problem, in which only tidal forces, and not the orbital motion and rotation, are taken into account [136]. Finally, Aizenman [2] investigated the so-called Roche-Riemann problem, in which internal fluid motions of uniform vorticity are considered. In these cases, the satellite star is not tidally locked to the primary. (It turns out this situation applies to neutron stars, in which the viscosity is too low to force corotation [33, 141].)

A fairly complete description of these problems is given in Chandrasekhar's book *Ellipsoidal Figures of Equilibrium* [50]. A virial method is used to find exact solutions for the equilibrium configurations. For example, consider the basic Roche problem. Define m_p as the mass of the primary star (producer of the tidal field), m_s as the mass of the distorted secondary, q as the mass ratio m_p/m_s, and $r_{\rm orb}$ as the orbital radius. Then $\psi \equiv Gm_p/r_{\rm orb}^3$ is the strength of the Newtonian tidal field of the primary star, and $\Omega^2 = G(m_p + m_s)/r_{\rm orb}^3 = (1 + 1/q)\psi$ is the square of the Keplerian orbital frequency of the orbit. Chandrasekhar shows that the principal axes of the secondary, a_1, a_2, and a_3 (where $a_1 > a_2 > a_3$), are given by the equations

$$\left[\left(3 + \frac{1}{q}\right) a_1^2 + a_3^2\right] \frac{\psi}{\pi\rho} = 2(A_1 a_1^2 - A_3 a_3^2)\,, \tag{4.1}$$

$$\left(\frac{1}{q} a_2^2 + a_3^2\right) \frac{\psi}{\pi\rho} = 2(A_2 a_2^2 - A_3 a_3^2)\,, \tag{4.2}$$

where $\rho = 3m_s/4\pi R^3$ is the average density of the secondary, R is its radius, and

$$A_i = a_1 a_2 a_3 \int_0^\infty \frac{du}{(a_1^2 + u)(a_2^2 + u)(a_3^2 + u)(a_i^2 + u)}\,. \tag{4.3}$$

This integral can be evaluated to get

$$A_1 = \frac{2a_2 a_3}{a_1^2 \sin^3\phi \sin^2\theta}[F(\theta,\phi) - E(\theta,\phi)]\,, \tag{4.4}$$

$$A_2 = \frac{2a_2 a_3}{a_1^2 \sin^3\phi \sin^2\theta \cos^2\theta}\left[E(\theta,\phi) - F(\theta,\phi)\cos^2\theta - \frac{a_3}{a_2}\sin^2\theta \sin\phi\right]\,, \tag{4.5}$$

$$A_3 = \frac{2a_2 a_3}{a_1^2 \sin^3\phi \cos^2\theta}\left[\frac{a_2}{a_3}\sin\phi - E(\theta,\phi)\right]\,, \tag{4.6}$$

where

$$E(\theta,\phi) = \int_0^\phi (1-\sin^2\theta\sin^2\phi)^{1/2} d\phi\,, \tag{4.7}$$

$$F(\theta,\phi) = \int_0^\phi (1-\sin^2\theta\sin^2\phi)^{-1/2} d\phi\,, \tag{4.8}$$

and

$$\sin\theta = \left(\frac{a_1^2-a_2^2}{a_1^2-a_3^2}\right)^{1/2}, \tag{4.9}$$

$$\cos\phi = \frac{a_3}{a_1}\,. \tag{4.10}$$

In his book, Chandrasekhar gives a table of a_2/a_1 and a_3/a_1 versus Ω^2 for the two limits $q=\infty$ (test mass) and $q=1$ (equal mass).

The Jeans problem is somewhat simpler, in that the rotation of the star is ignored. In this sense, it can be considered a "pure" tidal problem. Physically, we can imagine the secondary star freely falling on a radial trajectory toward the primary. We then move into a frame that is moving with the star. Chandrasekhar shows that the solutions are the same as those of the Roche problem, (4.1) and (4.2), but with $q=-1$ (dropping the interpretation of q as the mass ratio of the binary). In addition, $a_2 = a_3 < a_1$, so the equilibrium figures are prolate spheroids. In this case, the expressions for the A_i can be simplified. Defining the eccentricity

$$e = \left(1-\frac{a_2^2}{a_1^2}\right)^{1/2}, \tag{4.11}$$

we find that

$$A_1 = \frac{1-e^2}{e^3}\ln\frac{1+e}{1-e} - 2\frac{1-e^2}{e^2}\,, \tag{4.12}$$

$$A_2 = A_3 = \frac{1}{e^2} - \frac{1-e^2}{2e^3}\ln\frac{1+e}{1-e}\,. \tag{4.13}$$

The tidal field can then be equated to the eccentricity using

$$\frac{\psi}{\pi\rho} = \frac{1-e^2}{e^3}\ln\frac{1+e}{1-e} - \frac{6}{e^2}\frac{1-e^2}{3-e^2}\,. \tag{4.14}$$

A number of later authors [91, 172, 119, 222] solved the *relativistic* Jeans, Roche, and Roche-Riemann problems in which a star (usually a white dwarf or neutron star) is placed in the relativistic tidal field of a Schwarzschild or Kerr black hole. We derive this tidal field below.

An important limitation of the work by Chandrasekhar and these other authors is that they all assume an *incompressible* fluid. In the language of polytropes, in which the pressure is a power law function of density, $P = K_p\rho^\gamma \equiv K_p\rho^{(1+1/n)}$, an incompressible fluid corresponds to $\gamma = \infty$ or $n = 0$. This produces a constant density star. The extension of Chandrasekhar's results to compressible fluids was done by Lai, Rasio, and Shapiro [152, 153]. Like Chandrasekhar, they assume that surfaces of constant density are self-similar ellipsoids. While this is true exactly in the incompressible case, it is only an approximation in the compressible case. Wiggins and Lai [254] used these results to solve the relativistic Roche-Riemann problem for a compressible star.

This work has, for several years, been the standard for the study of tidal disruption. However, it ignores two very important features of the problem. The first relates to the equation of state. As mentioned above, the equation of state of a neutron star is unknown and is, in fact, a primary reason for doing such calculations in the first place. However, the equation of state is almost certainly not described by a polytrope. It is necessary to investigate more realistic equations of state and their impact on the tidal disruption problem. Second, both the undistorted fluid star and its tidally perturbed structure are analyzed using Newtonian gravity. While this may be fine for a white dwarf, it is not appropriate for a neutron star, in which $m_{\rm NS}/R \sim 0.1 - 0.2$.

The first goal of this chapter, therefore, is to develop a relativistic version of the work of Wiggins and Lai. The approach we use is perturbation theory.

4.3 Metric of a perturbed neutron star

4.3.1 Unperturbed neutron star

The metric for a perturbed neutron star can be written as

$$g_{\mu\nu} = \hat{g}_{\mu\nu} + h_{\mu\nu}\,, \tag{4.15}$$

where $\hat{g}_{\mu\nu}$ is the unperturbed metric. (For clarity, we will always represent unperturbed quantities with hats.) The first step is to find $\hat{g}_{\mu\nu}$. The general spherically symmetric metric, expressed using Schwarzschild coordinates (t, r, θ, ϕ), is

$$\hat{g}_{\mu\nu} = \mathrm{diag}(-e^{2\Phi(r)}, e^{2\Lambda(r)}, r^2, r^2\sin^2\theta)\,. \tag{4.16}$$

The forms of the functions $\Phi(r)$ and $\Lambda(r)$ are determined by demanding that the metric satisfy the Einstein equations $G_{\mu\nu} = 8\pi T_{\mu\nu}$. When there is no matter present, as for a black hole or in the exterior of a star, we recover the familiar Schwarzschild solution (1.22). When there is matter present, we also need to solve for the unperturbed pressure $\hat{P}(r)$ and the unperturbed total energy density $\hat{\rho}(r)$. The two are related by the equation of state, $P(\rho)$.

We assume that the matter can be treated as a perfect fluid, which has no shear stress, viscosity, or heat conduction. The stress-energy tensor of a perfect fluid can be written $T_{\mu\nu} = (\rho + P)u_\mu u_\nu + P g_{\mu\nu}$, where u^μ is the four-velocity of the fluid. In the static case, the four-velocity has only a time component. To normalize it, we require that $g_{\mu\nu}u^\mu u^\nu = -1$, giving at zeroth order

$$\hat{u}^\mu = (e^{-\Phi(r)}, 0, 0, 0) \tag{4.17}$$

and

$$\hat{u}_\mu = (-e^{\Phi(r)}, 0, 0, 0)\,. \tag{4.18}$$

The unperturbed stress-energy tensor is then

$$\hat{T}_{\mu\nu} = \mathrm{diag}(\hat{\rho}e^{2\Phi(r)}, \hat{P}e^{2\Lambda(r)}, \hat{P}r^2, \hat{P}r^2\sin^2\theta)\,. \tag{4.19}$$

Writing out the Einstein equations, we find that $\hat{G}_{tt} = 8\pi\hat{T}_{tt}$ gives

$$\frac{e^{-2\Lambda}}{r^2}(2r\Lambda' - 1 + e^{2\Lambda}) = 8\pi\hat{\rho}\,, \tag{4.20}$$

where $'$ signifies radial derivatives d/dr. This equation will be useful in its own right as an expression for Λ' but can also be rewritten in a more physical form. If we define an enclosed mass,

$$m(r) = \frac{r}{2}(1 - e^{-2\Lambda})\,, \tag{4.21}$$

then

$$m' = 4\pi r^2\hat{\rho}\,, \tag{4.22}$$

analogous to the Newtonian mass equation, except that $\hat{\rho}$ now represents the total energy density, not just rest mass density. At the surface of the star ($r = R$), $m(R) = m_{\mathrm{NS}}$, the mass of the neutron star and the mass which appears in the exterior metric (1.22). Moving on, $\hat{G}_{rr} = 8\pi\hat{T}_{rr}$ gives

$$\frac{e^{-2\Lambda}}{r^2}(2r\Phi' + 1 - e^{2\Lambda}) = 8\pi\hat{P}\,, \tag{4.23}$$

which with the definition of m can be rewritten as

$$\Phi' = \frac{m + 4\pi r^3\hat{P}}{r(r - 2m)}\,. \tag{4.24}$$

Finally, the conservation of the stress-energy (specifically $\nabla_\mu\hat{T}^{\mu r} = 0$) can be used to obtain a third equation:

$$\hat{P}' = -(\hat{\rho} + \hat{P})\Phi'\,. \tag{4.25}$$

These are the famous Tolman-Oppenheimer-Volkoff (TOV) equations [244, 184].

4.3.2 General perturbations to a spherically symmetric spacetime

The problem of perturbations to a spherically symmetric spacetime was first studied by Regge and Wheeler in vacuum [204] and extended to fluid stars by Thorne and Campolattaro [239]. Our treatment will follow theirs very closely. However, we will use a more modern notation and signature for the unperturbed metric, which affects signs and factors of 2 throughout the calculation.

All perturbed quantities can be expanded in spherical harmonics. Quantities which transform as scalars under rotations are expanded in the familiar scalar spherical harmonics $Y_{\ell m}(\theta, \phi)$, which have parity $(-1)^\ell$. Regge and Wheeler call this "even parity"; it can also be referred to as "polar." Vector quantities can be expanded in vector spherical harmonics, which have an even, or polar type,

$$\Psi_{\ell m j} = \partial_j Y_{\ell m}(\theta, \phi)\,, \tag{4.26}$$

and an "odd," or "axial" type (with parity $(-1)^{\ell+1}$),

$$\Phi_{\ell m j} = \epsilon_j{}^k \partial_k Y_{\ell m}(\theta, \phi)\,, \tag{4.27}$$

where

$$\epsilon_\theta{}^\phi = -\frac{1}{\sin\theta}\,, \tag{4.28}$$

$$\epsilon_\phi{}^\theta = \sin\theta\,, \tag{4.29}$$

$$\epsilon_\theta{}^\theta = \epsilon_\phi^\phi = 0\,. \tag{4.30}$$

(Note that in this section, Latin indices range only over the angles θ and ϕ.) Finally, tensor quantities are expanded in terms of tensor spherical harmonics, of even parity,

$$\Psi_{\ell m jk} = Y_{\ell m|jk}\,, \tag{4.31}$$

$$\Phi_{\ell m jk} = \gamma_{jk} Y_{\ell m}\,, \tag{4.32}$$

	Odd parity	Even parity
ξ_r	0	$X(r,t)Y_{\ell m}$
ξ_θ	$U(r,t)\Phi_{\ell m\theta}$	$V(r,t)\Psi_{\ell m\theta}$
ξ_ϕ	$U(r,t)\Phi_{\ell m\phi}$	$V(r,t)\Psi_{\ell m\phi}$
h_{tt}	0	$e^{2\Phi(r)}H_0(r,t)Y_{\ell m}$
h_{tr}	0	$H_1(r,t)Y_{\ell m}$
h_{rr}	0	$e^{2\Lambda(r)}H_2(r,t)Y_{\ell m}$
$h_{tj}, j \in \{\theta,\phi\}$	$h_0(r,t)\Phi_{\ell mj}$	$H_3(r,t)\Psi_{\ell mj}$
$h_{rj}, j \in \{\theta,\phi\}$	$h_1(r,t)\Phi_{\ell mj}$	$H_4(r,t)\Psi_{\ell mj}$
$h_{jk}, \{j,k\} \in \{\theta,\phi\}$	$h_2(r,t)\chi_{\ell mjk}$	$r^2G(r,t)\Psi_{\ell mjk} + r^2K(r,t)\Phi_{\ell mjk}$

Table 4.1: General form of nonradial perturbations, split into odd and even parity and expanded as scalar, vector, and tensor spherical harmonics.

and odd parity,

$$\chi_{\ell mjk} = \frac{1}{2}(\epsilon_j{}^n\Psi_{\ell mnk} + \epsilon_k{}^n\Psi_{\ell mnj})\,, \tag{4.33}$$

where γ_{jk} is the metric

$$\gamma_{\theta\theta} = 1\,, \tag{4.34}$$

$$\gamma_{\theta\phi} = \gamma_{\phi\theta} = 0\,, \tag{4.35}$$

$$\gamma_{\phi\phi} = \sin^2\theta\,, \tag{4.36}$$

and the symbol $|$ means to take a covariant derivative with respect to this metric.

The quantities we want to expand into spherical harmonics are the 10 independent components of the metric perturbation $h_{\mu\nu}$ and the three components of the fluid displacement $\xi = (\xi^r, \xi^\theta, \xi^\phi)$. ξ^r, h_{tt}, h_{tr}, and h_{rr} behave like scalars under rotation, while $(\xi^\theta, \xi^\phi), (h_{t\theta}, h_{t\phi})$, and $(h_{r\theta}, h_{r\phi})$ are vectors. The remainder of h transforms as a tensor. We can then write the most general odd and even parity perturbations; see Table 4.1.

Several simplifications can be made to these perturbations. First, we can pick a specific gauge. A useful gauge is that used by Regge and Wheeler [204], which is chosen to eliminate $h_2(r,t)$ in the odd parity case and $H_3(r,t), H_4(r,t)$, and $G(r,t)$ in the even parity case. Next, we can (temporarily) set $m = 0$. This choice has no effect

when solving for the radial perturbations and makes that calculation simpler. We will later restore the proper angular dependence when it is needed. Third, we redefine the fluid displacement coefficients to match a convention of Thorne and Campolattaro [239]. Finally, we eliminate all time dependence from the perturbation functions and look only at the *static* perturbations to the star. (This is similar to the Roche, Jeans, Darwin, and Roche-Riemann problems, which ignore dynamics and focus instead on the existence of equilibrium solutions.)

It turns out that we will only be interested in the even parity, polar perturbations. The h_{rr} term, not present in the odd perturbations, is the crucial one for calculating the principal axes of the perturbed star. When we later change focus to investigate Love numbers, h_{tt} will be the critical quantity.[3] The simplified even metric perturbation is

$$h_{\mu\nu} = \begin{bmatrix} e^{2\Phi}H_0(r) & H_1(r) & 0 & 0 \\ H_1(r) & e^{2\Lambda}H_2(r) & 0 & 0 \\ 0 & 0 & r^2K(r) & 0 \\ 0 & 0 & 0 & r^2\sin^2\theta K(r) \end{bmatrix} P_\ell(\cos\theta)\,, \tag{4.37}$$

where $P_\ell(x)$ are the Legendre polynomials of index ℓ. (For now, we have absorbed some normalization constants into H_0, H_1, H_2, and K. When we later restore $m \neq 0$, we will pull these back out into the full spherical harmonic $Y_{\ell m}$.) The fluid displacement vector is (with the aforementioned redefinition)

$$\xi^r = r^{-2}e^{-\Lambda}W(r)P_\ell(\cos\theta)\,, \tag{4.38}$$

$$\xi^\theta = -r^{-2}V(r)\partial_\theta P_\ell(\cos\theta)\,, \tag{4.39}$$

$$\xi^\phi = 0\,. \tag{4.40}$$

[3]The effect of the odd parity, axial perturbations is to cause a differential rotation [239].

4.3.3 Even parity perturbation equations

In the static case, the fluid displacement functions W and V are independent of time, so there are no modifications to the r, θ, and ϕ components of four-velocity. There will, however, be a modification to the time component:

$$\hat{g}_{tt}\hat{u}^t\hat{u}^t + 2\hat{g}_{tt}\hat{u}^t\delta u^t + h_{tt}\hat{u}^t\hat{u}^t = -1 \Rightarrow \delta u^t = \frac{1}{2}e^{-\Phi}H_0(r)P_\ell(\cos\theta)\,. \tag{4.41}$$

The perturbed stress-energy tensor is, in general, given by

$$\delta T_{\mu\nu} = (\delta\rho + \delta P)\hat{u}_\mu\hat{u}_\nu + (\hat{\rho} + \hat{P})\delta u_\mu\hat{u}_\nu + (\hat{\rho} + \hat{P})\hat{u}_\mu\delta u_\nu + \delta P\hat{g}_{\mu\nu} + \hat{P}h_{\mu\nu}\,, \tag{4.42}$$

where

$$\delta u_\mu = \hat{g}_{\mu\nu}\delta u^\nu + h_{\mu\nu}\hat{u}^\nu\,. \tag{4.43}$$

It turns out to be convenient to calculate everything with mixed indices, in which case the stress-energy tensor becomes[4]:

$$\delta T_\mu{}^\nu = \begin{bmatrix} -\delta\rho & 0 & 0 & 0 \\ e^{-2\Phi}H_1(\hat{\rho} + \hat{P})P_\ell(\cos\theta) & \delta P & 0 & 0 \\ 0 & 0 & \delta P & 0 \\ 0 & 0 & 0 & \delta P \end{bmatrix}. \tag{4.44}$$

In [239], $\delta\rho$ and δP are expressed in terms of W and V; however, since W and V are time independent and do not appear in the modified four-velocity, it is simpler to just ignore the fluid displacement and use $\delta\rho$ and δP as the variables of interest. Further, $\delta\rho = \delta P/(dP/d\rho)$ given an equation of state $P(\rho)$, so we can focus only on δP.

We now calculate the perturbed Einstein tensor. First, the perturbed Christoffel symbols are given by

$$\delta\Gamma^\sigma{}_{\mu\nu} = \frac{1}{2}\hat{g}^{\sigma\rho}(\hat{\nabla}_\mu h_{\rho\nu} + \hat{\nabla}_\nu h_{\mu\rho} - \hat{\nabla}_\rho h_{\mu\nu})\,, \tag{4.45}$$

[4]The symmetry of the tensor is hidden because we are using mixed indices.

where $\hat{\nabla}$ means a covariant derivative taken with respect to the unperturbed metric $\hat{g}_{\mu\nu}$. The perturbed Ricci tensor is given by

$$\delta R_{\mu\nu} = \hat{\nabla}_\lambda \delta\Gamma^\lambda_{\mu\nu} - \hat{\nabla}_\nu \delta\Gamma^\lambda_{\mu\lambda}\,. \tag{4.46}$$

Finally, the perturbed Einstein tensor is

$$\delta G_{\mu\nu} = \delta R_{\mu\nu} - \frac{1}{2}\hat{g}_{\mu\nu}\delta R - \frac{1}{2}h_{\mu\nu}\hat{R}\,, \tag{4.47}$$

where $\hat{R} = \hat{g}^{\mu\nu}\hat{R}_{\mu\nu}$ is the unperturbed Ricci scalar and

$$\delta R = -h^{\mu\nu}\hat{R}_{\mu\nu} + \hat{g}^{\mu\nu}\delta R_{\mu\nu} \tag{4.48}$$

is the perturbed version. We now set $\delta G_\mu{}^\nu = 8\pi\delta T_\mu{}^\nu$. To begin, $\delta G_\theta{}^\theta - \delta G_\phi{}^\phi = 8\pi(\delta T_\theta{}^\theta - \delta T_\phi{}^\phi) = 0$ tells us that $H_0 = H_2$. This fact, which is true even in the nonstatic case, allows us to eliminate H_2 from all equations. Next, $\delta G_t{}^r = 8\pi\delta T_t{}^r = 0$ and $\delta G_r{}^t = 8\pi\delta T_r{}^t$ both give $H_1 = 0$. This result is specific to the static case. $\delta G_t{}^\theta$ and $\delta G_\theta{}^t$ both contain only terms involving H_1, so they are irrelevant in the static case.

The first substantial relation comes from $\delta G_r{}^\theta = 8\pi\delta T_r{}^\theta = 0$ (or $\delta G_\theta{}^r = 8\pi\delta T_\theta{}^r = 0$):

$$H_0' - K' + 2\Phi' H_0 = 0\,. \tag{4.49}$$

$\delta G_r{}^r = 8\pi\delta T_r{}^r$ gives

$$\begin{aligned}&\frac{1}{2r^2}e^{-2\Lambda}[-e^{2\Lambda}(\ell(\ell+1)-2)K + (-2+e^{2\Lambda}\ell(\ell+1) - 4r\Phi')H_0 - 2rH_0' \\ &+ 2r(1+r\Phi')K']P_\ell(\cos\theta) = 8\pi\delta P\,.\end{aligned} \tag{4.50}$$

To find δP, the simplest method is to take derivatives of the stress-energy tensor, which is locally conserved. $\delta(\nabla_\mu T_\theta{}^\mu) = 0$ gives

$$\delta P = \frac{1}{2}(\hat{\rho} + \hat{P})H_0 P_\ell(\cos\theta)\,. \tag{4.51}$$

Of the other three components, one is zero, one gives $H_1 = 0$, and the third (the r equation) can be derived from the θ equation. Similarly, the tt, $\theta\theta$, and $\phi\phi$ components of the Einstein equation can be derived from the previous equations. All other components of the Einstein equation are zero. We can solve (4.49), (4.50), and (4.51) for H_0' and K'. The results are:

$$K' = \frac{1}{\Phi'}\left(\frac{1}{r^2} - \frac{e^{2\Lambda}\ell(\ell+1)}{2r^2} + 4\pi e^{2\Lambda}(\hat{\rho}+\hat{P})\right)H_0 + \frac{1}{\Phi'}\left(\frac{\ell(\ell+1)}{2r^2} - \frac{1}{r^2}\right)e^{2\Lambda}K\,, \tag{4.52}$$

$$H_0' = K' - 2\Phi' H_0\,. \tag{4.53}$$

As we shall see later, we are interested only in $\ell = 2$. In this case, (4.52) and (4.53) become

$$K' = \frac{1}{\Phi'}\left(\frac{1}{r^2} - \frac{3e^{2\Lambda}}{r^2} + 4\pi e^{2\Lambda}(\hat{\rho}+\hat{P})\right)H_0 + \frac{2e^{2\Lambda}}{\Phi' r^2}K\,, \tag{4.54}$$

$$H_0' = K' - 2\Phi' H_0\,. \tag{4.55}$$

These are the central equations of this chapter.

4.3.4 Boundary conditions: origin

We only want to consider solutions which are well behaved at the origin. Such solutions can be expanded in a Taylor series at $r = 0$. These expressions can then be used to start a numerical integration of (4.54) and (4.55). It is useful to first write down the expansions of unperturbed quantities at small r. To the orders we will need, we find

$$m(r) = \frac{4}{3}\pi\hat{\rho}_c r^3 - \frac{8\pi^2}{15}\frac{(\hat{\rho}_c+\hat{P}_c)(\hat{\rho}_c+3\hat{P}_c)}{(dP/d\rho)_c}r^5\,, \tag{4.56}$$

$$\hat{P}(r) = \hat{P}_c - \frac{2\pi}{3}(\hat{\rho}_c+\hat{P}_c)(\hat{\rho}_c+3\hat{P}_c)r^2\,, \tag{4.57}$$

$$\hat{\rho}(r) = \hat{\rho}_c - \frac{2\pi}{3}\frac{(\hat{\rho}_c + \hat{P}_c)(\hat{\rho}_c + 3\hat{P}_c)}{(dP/d\rho)_c} r^2 \,, \tag{4.58}$$

$$\Phi(r) = \Phi_c + \frac{2\pi}{3}(\hat{\rho}_c + 3\hat{P}_c) r^2 - \frac{2\pi^2}{45}(\hat{\rho}_c + 3\hat{P}_c)\left(15\hat{P}_c - 5\hat{\rho}_c + \frac{3(\hat{\rho}_c + \hat{P}_c)}{(dP/d\rho)_c}\right) r^4 \,, \tag{4.59}$$

where $\hat{P}_c$, $\hat{\rho}_c$, Φ_c, and $(dP/d\rho)_c$ are the values at $r = 0$. Now we do the same for the perturbations. Specializing again to $\ell = 2$, we find

$$H_0(r) = c_0 r^2 \left[1 - \frac{2\pi}{7}\left(\frac{1}{3}\hat{\rho}_0 + 11\hat{P}_0 + \frac{\hat{\rho}_0 + \hat{P}_0}{(dP/d\rho)_0}\right) r^2\right] \,, \tag{4.60}$$

$$K(r) = c_0 r^2 \left[1 - \frac{2\pi}{7}\left(-2\hat{\rho}_0 + 4\hat{P}_0 + \frac{\hat{\rho}_0 + \hat{P}_0}{(dP/d\rho)_0}\right) r^2\right] \,. \tag{4.61}$$

4.3.5 Boundary conditions: large r

Outside of the star, (4.52) and (4.53) simplify to

$$K' = \left(\frac{r - 2m_{\rm NS}}{r m_{\rm NS}} - \frac{\ell(\ell+1)}{2m_{\rm NS}}\right) H_0 + \frac{1}{m_{\rm NS}}\left(\frac{\ell(\ell+1)}{2} - 1\right) K \,, \tag{4.62}$$

$$H_0' = K' - \frac{2m_{\rm NS}}{r(r - 2m_{\rm NS})} H_0 \,. \tag{4.63}$$

We can find analytic solutions to these equations. To do so, it is convenient to write them as a second order differential equation:

$$H_0'' + \left(\frac{2}{r} + \frac{2m_{\rm NS}}{r(r - 2m_{\rm NS})}\right) H_0' - \left(\frac{\ell(\ell+1)}{r(r - 2m_{\rm NS})} + \frac{4m_{\rm NS}^2}{r^2(r - 2m_{\rm NS})^2}\right) H_0 = 0 \,. \tag{4.64}$$

Changing variables to $x = r/m_{\rm NS} - 1$ [108, 116], we find

$$(x^2 - 1)H_0'' + 2xH_0' - \left(\ell(\ell+1) + \frac{4}{x^2+1}\right) H_0 = 0 \,. \tag{4.65}$$

This is the associated Legendre differential equation. The solution is

$$H_0(r) = c_1 P_\ell^2\left(\frac{r}{m_{\rm NS}} - 1\right) + c_2 Q_\ell^2\left(\frac{r}{m_{\rm NS}} - 1\right) \,, \tag{4.66}$$

where P_ℓ^m and Q_ℓ^m are the associated Legendre functions of the first and second kind. In the case $\ell = 2$, (4.66) is

$$\begin{aligned} H_0(r) = c_1 &\left(\frac{r}{m_{\rm NS}}\right)^2 \left(1 - \frac{2m_{\rm NS}}{r}\right) \\ &+ c_2 \left(\frac{r}{m_{\rm NS}}\right)^2 \left(1 - \frac{2m_{\rm NS}}{r}\right) \left[-\frac{m_{\rm NS}(m_{\rm NS} - r)(2m_{\rm NS}^2 + 6m_{\rm NS}r - 3r^2)}{r^2(2m_{\rm NS} - r)^2}\right. \\ &\left. -\frac{3}{2}\ln\left(1 - \frac{2m_{\rm NS}}{r}\right)\right] . \end{aligned} \tag{4.67}$$

We can also solve for K:

$$\begin{aligned} K(r) = c_1 &\left(\frac{r}{m_{\rm NS}}\right)^2 \left(1 - \frac{2m_{\rm NS}^2}{r^2}\right) \\ &+ c_2 \left[-3 + \frac{2m_{\rm NS}}{r} - \frac{3r}{m_{\rm NS}} + \left(3 - \frac{3}{2}\frac{r^2}{m_{\rm NS}^2}\right)\ln\left(1 - \frac{2m_{\rm NS}}{r}\right)\right] . \end{aligned} \tag{4.68}$$

It is not obvious, but in each case, the second solution dies away at large r, as r^{-3}. In Sec. 4.4, we will use these analytic exterior solutions to match the tidal field at $r \to \infty$ and to match the interior solution at $r = R_+$.

4.3.6 Boundary conditions: surface

In addition to boundary conditions at small and large r, we also need boundary conditions at the surface of the star. We can also call these "junction conditions," since they tell how to join the perturbations inside the star to the analytic exterior solutions derived above. The general conditions for the junction of two metrics at a hypersurface separating them were derived by Darmois [65]. Let n^μ be the normal vector to the surface. In our case, the surface is timelike, so n^μ is spacelike, and $n_\mu n^\mu = 1$. If the stress-energy is not singular on the surface (e.g. a "surface layer"), then two quantities are continuous across the surface. The first is the *induced metric*, or first fundamental form,

$$i_{\mu\nu} \equiv g_{\mu\nu} - n_\mu n_\nu \,, \tag{4.69}$$

and the other is the *extrinsic curvature*[5], or second fundamental form,

$$e_{\mu\nu} \equiv (\nabla_\alpha n_\mu) i^\alpha{}_\nu = \nabla_\nu n_\mu - (\nabla_\alpha n_\mu) n^\alpha n_\nu \,. \tag{4.70}$$

In the perturbative case, an interesting question is whether the upper or lower components of these tensors are continuous. The question is answered by Martín-García and Gundlach (MG) [171], whose work is the basis of this section. Let the surface be defined by the level surface $f = 0$ of a scalar field f. The normal vector can be defined by the gradient of the field:

$$n_\mu = \frac{\partial_\mu f}{(g^{\lambda\rho}\partial_\lambda f \partial_\rho f)^{1/2}} \,. \tag{4.71}$$

For the surface of a star, $f = -P$, where the negative sign is necessary for the normal vector defined above to point the correct way. Martín-García and Gundlach argue that tensors which are intrinsic to the hypersurface can be contracted on an index with $n_\mu \propto \partial_\mu f$ to get zero. This means that contravariant tensors, with superscript indices, are truly intrinsic, while covariant tensors, with subscript indices, are only intrinsic in a metric-dependent way. Following this argument, MG proceed to calculate $\delta i^{\mu\nu}$ and $\delta e^{\mu\nu}$ and force them to be continuous across the surface.

However, there is another complication. The surface of the star is defined physically as the location at which the pressure drops to zero. In the unperturbed system, the surface is defined to be located at $r = R$. However, the location of the surface in the perturbed system depends on the choice of gauge. In general, the unperturbed surface and perturbed surface will *not* line up. However, it is much easier to analyze the system in the case that they do, so MG perform their matching in this special "surface gauge." They then transform back to a general gauge.

[5]The extrinsic curvature is also a key quantity in numerical relativity. In that case, the surfaces in question are spacelike slices of spacetime, and the normal vector is timelike.

The final junction conditions in Regge-Wheeler gauge are:

$$[H_0] = [K] = 0\,, \tag{4.72}$$

$$[H_0'] = [K'] = \frac{4\pi e^{2\Lambda}}{\Phi'}[\hat{\rho}]H_0\,, \tag{4.73}$$

where brackets denote the change in a quantity across the boundary. We see that the junction conditions are trivial (i.e., the perturbations and their derivatives are continuous across the boundary) unless the density is discontinuous. The same junction conditions will thus apply at any surface *within* the star in which the density changes discontinuously. This occurs for some of the equations of state that we describe later. Inspection of (4.54) and (4.55) shows that the continuity of the unperturbed quantities and (4.72) imply (4.73). The junction conditions are thus "built into" the differential equations, so we do not need to treat any of these cases very carefully.

Price and Thorne [202], following up the work of Thorne and Campolattaro, found the result

$$[H_0'] = [K'] = 8\pi r^{-2} e^{\Lambda}[\hat{\rho}]W\,, \tag{4.74}$$

with W as defined in (4.39) and the signs and factors of 2 adjusted to our convention. Comparing to (4.73), we find that

$$W = \frac{r^2 e^{\Lambda} H_0}{2\Phi'}\,, \tag{4.75}$$

so that the surface is displaced in Regge-Wheeler gauge by

$$\xi^r = \frac{H_0}{2\Phi'} P_\ell(\cos\theta) = -\frac{\delta P}{\partial_r P}\,. \tag{4.76}$$

This matches our intuition that the *Lagrangian* change of pressure

$$\Delta P \equiv \delta P + (\partial_r P)\xi^r = 0 \tag{4.77}$$

at the surface (and, in fact, throughout the star).

4.4 Metric of a neutron star in a binary

So far we have investigated the nonradial perturbations of a neutron star, focusing on the static $\ell = 2$ perturbations in Regge-Wheeler gauge. However, we have not dealt with the actual physical situation in which we are interested: the structure of a tidally distorted neutron star in a black hole-neutron star binary. In this section, we will describe a method to analyze this scenario. In the end, we will discover that the static $\ell = 2$ Regge-Wheeler perturbations are exactly what we need.

To construct the metric of the distorted neutron star, we follow the method of Alvi [8]. Alvi computes the metric of a black hole binary including tidal effects on the black holes, with the goal of providing a more accurate set of initial data for numerical relativity simulations. He begins by dividing spacetime into four regions. Define m_1 and m_2 to be the masses of the holes, d to be their separation in harmonic coordinates, r_1 and r_2 to be the distances from each hole, and r_{CM} to be the distance from the center of mass. The four regions are:

- Region I: A region local to hole 1, $r_1 < r_1^{in} = \sqrt{m_1 d}$
- Region II: A region local to hole 2, $r_2 < r_2^{in} = \sqrt{m_2 d}$
- Region III: The rest of the *near zone*, excluding the previous two regions, i.e. $r_1 > r_1^{in}, r_2 > r_2^{in}, r_{CM} < r^{out} = \lambda_c/2\pi = \sqrt{d^3/(m_1+m_2)}/2$
- Region IV: The *radiation zone*, $r_{CM} > r^{out}$

Here $\lambda_c/2\pi$ is the characteristic wavelength of gravitational radiation emitted by the binary: $\lambda_c = \pi/\Omega$. The metric in regions I and II is the Schwarzschild solution (since Alvi restricts the black holes to be nonspinning), modified by a perturbation due to the tidal interaction. The metrics in regions III and IV are constructed by standard post-Newtonian techniques.

There also exist so-called "buffer" or overlap zones between pairs of touching regions (i.e., I and III, II and III, III and IV). In these buffer zones, *both* metrics are valid. The metrics of the perturbed holes in regions I and II must be properly matched to the post-Newtonian metric of region III. (The metrics of III and IV, taken from earlier post-Newtonian studies, already match in their buffer region by construction.) By performing this match, we can determine the unknown coefficients of the perturbed hole metrics as well as the transformation between the coordinate systems in each region. This procedure is known as "matched asymptotic expansion" [240]. (It should be noted that Alvi actually does this matching incorrectly [261]).

In practice, Alvi takes a slightly different approach. Instead of doing the full matching procedure at once, he first calculates the electric and magnetic tidal fields of one hole (say, hole 2) in the buffer zone of the other (hole 1). In doing this, he explicitly treats hole 1 as a test mass on a geodesic orbit around hole 2. A simpler match determines the unknown coefficients of the region I metric. The complete matching procedure is then used to relate the coordinate systems of regions I and III. (The whole procedure is, of course, repeated for the tidal field of hole 1 in the buffer zone of hole 2.) The disadvantage of this system is that it requires some *ad hoc* fixes due to the finite mass of hole 1. It also only correctly produces the leading-order tidal fields within the post-Newtonian expansion [193]. A better approach [236] is to perform a complete matched asymptotic expansion between the two regions.

It is easy to imagine extending Alvi's procedure to include a neutron star. Outside the star, the metric is the same as for a black hole. Therefore, in the matching region, the problem is exactly the same. As we are only interested in determining the structure of the tidally perturbed star, we can ignore the details of the post-Newtonian metric and the full asymptotic matching procedure. In this case, therefore, Alvi's simplified matching method is useful, even if somewhat incorrect.

4.4.1 Tidal field of the black hole

Alvi's matching procedure is carried out in the local asymptotic rest frame (LARF) of the neutron star. The LARF, as defined by Thorne and Hartle [240], is "a coordinate

system in the buffer region that is as nearly globally inertial and Lorentz as possible and in which the body is momentarily at rest." In the buffer zone, $R \ll r \ll \mathcal{L}$, where R is the radius of the neutron star, r is the distance from the star, and $\mathcal{L}$ is the length scale on which the black hole's tidal field changes.[6] The metric can thus be expanded as follows:

$$g = g^{(0)} + m_{\rm NS} g^{(1)} + m_{\rm NS}^2 g^{(2)} + \cdots . \tag{4.78}$$

Here $g^{(0)}$ is the external universe without the neutron star (i.e., the black hole) and the other terms are perturbations due to the presence of the neutron star, expanded in powers of the star's mass $m_{\rm NS}$. The goal is to obtain an expression for $g^{(0)}$ in LARF coordinates. As stated above, we can do this by considering a test observer who freely falls along a geodesic of the black hole geometry. Later, we can add effects due to the finite mass of the neutron star. We learned in Chapter 1 that a freely falling frame is *locally* inertial; however, it deviates from flat space as the distance from the geodesic r increases. To lowest order in r, the metric is:

$$g_{00} = -1 - R_{0i0j}(t) x^i x^j + O(r^3) , \tag{4.79}$$

$$g_{0i} = -\frac{2}{3} R_{0jik}(t) x^j x^k + O(r^3) , \tag{4.80}$$

$$g_{ij} = \delta_{ij}[1 - R_{0k0m}(t) x^k x^m] + O(r^3) . \tag{4.81}$$

In this expression, the Riemann tensor is evaluated at the geodesic ($r = 0$). g_{ij} can also be expressed in terms of the *electric tidal field* $\mathcal{E}_{ij} = R_{0i0j}$ and the *magnetic tidal field* $\mathcal{B}_{ij} = \frac{1}{2}\epsilon_{ipq} R_{pqj0}$, where ϵ_{ipq} is the three-dimensional Levi-Civita symbol. As we shall see, electric tidal terms generate even parity perturbations, while magnetic tidal terms generate odd parity perturbations.

The particular coordinate choice here is not the familiar Fermi normal coordinates [169], but rather those used by Thorne and Hartle [240] and developed further by Zhang [262] (and called "THZ coordinates" by later authors). They are harmonic

[6] This approximation is actually quite poor for close orbits, one of many things which may affect our results.

coordinates; that is, $\Box x^\mu = 0$. The usefulness of these coordinates will become apparent later.

Marck [170] computed the components of an orthonormal basis (a *tetrad*) for the freely falling frame. Using the Marck tetrad, and simplifying to the Schwarzschild case, we can derive the components of the Riemann tensor in the freely falling frame. They are [8]:

$$R_{0101} = \frac{m_{\rm BH}}{d^3}\left[1 - 3\left(1 + \frac{\mathcal{K}}{d^2}\right)\cos^2\Omega t\right], \tag{4.82}$$

$$R_{0202} = \frac{m_{\rm BH}}{d^3}\left[1 - 3\left(1 + \frac{\mathcal{K}}{d^2}\right)\sin^2\Omega t\right], \tag{4.83}$$

$$R_{0303} = \frac{m_{\rm BH}}{d^3}\left(1 + \frac{3\mathcal{K}}{d^2}\right), \tag{4.84}$$

$$R_{0102} = R_{0201} = -\frac{3m_{\rm BH}}{d^3}\left(1 + \frac{\mathcal{K}}{d^2}\right)\cos\Omega t \sin\Omega t\,, \tag{4.85}$$

$$R_{0112} = -R_{0121} = R_{0323} = -R_{0332} = \frac{3m_{\rm BH}\sqrt{\mathcal{K}}}{d^4}\left(1 + \frac{\mathcal{K}}{d^2}\right)^{1/2}\cos\Omega t\,, \tag{4.86}$$

$$R_{0212} = -R_{0221} = R_{0331} = -R_{0313} = \frac{3m_{\rm BH}\sqrt{\mathcal{K}}}{d^4}\left(1 + \frac{\mathcal{K}}{d^2}\right)^{1/2}\sin\Omega t\,, \tag{4.87}$$

where $m_{\rm BH}$ is the black hole mass, d is the radius of the orbit,

$$\Omega = \sqrt{\frac{m_{\rm BH}}{d^3}} \tag{4.88}$$

is the rotation rate of the tidal field seen in the freely falling frame, and

$$\mathcal{K} = d^2\left(\frac{m_{\rm BH}}{d - 3m_{\rm BH}}\right) \tag{4.89}$$

is the relativistic part of the tide.[7] Plugging into (4.79)-(4.81), we find

$$g_{tt} = -1 + \frac{m_{\rm BH}}{d^3}\left[3\left(1 + \frac{\mathcal{K}}{d^2}\right)(x\cos\Omega t + y\sin\Omega t)^2 - (x^2 + y^2 + z^2) - \frac{3\mathcal{K}}{d^2}z^2\right], \tag{4.90}$$

[7] This is essentially the same way other authors compute the tide for the relativistic Jeans, Roche, and Roche-Riemann problems, except they express it directly as a force and not as part of the metric.

$$g_{tx} = \frac{2m_{\rm BH}\sqrt{\mathcal{K}}}{d^4}\left(1+\frac{\mathcal{K}}{d^2}\right)^{1/2}\left[(z^2-y^2)\sin\Omega t - xy\cos\Omega t\right], \tag{4.91}$$

$$g_{ty} = \frac{2m_{\rm BH}\sqrt{\mathcal{K}}}{d^4}\left(1+\frac{\mathcal{K}}{d^2}\right)^{1/2}\left[(x^2-z^2)\cos\Omega t + xy\sin\Omega t\right], \tag{4.92}$$

$$g_{tz} = \frac{2m_{\rm BH}\sqrt{\mathcal{K}}}{d^4}\left(1+\frac{\mathcal{K}}{d^2}\right)^{1/2}(y\cos\Omega t - x\sin\Omega t)z\,, \tag{4.93}$$

$$g_{ij} = \delta_{ij}\left(1+\frac{m_{\rm BH}}{d^3}\left[3\left(1+\frac{\mathcal{K}}{d^2}\right)(x\cos\Omega t + y\sin\Omega t)^2 - (x^2+y^2+z^2) - \frac{3\mathcal{K}}{d^2}z^2\right]\right). \tag{4.94}$$

At $t=0$, the x-axis is defined to be aligned with the black hole. The y-axis is in the direction of motion, and the z-axis is perpendicular to the plane of the orbit. Converting to spherical coordinates ($x = r\sin\theta\cos\phi, y = r\sin\theta\sin\phi, z = \cos\theta$), this can be written as the line element

$$\begin{aligned} ds^2 = &-dt^2 + dr^2 + r^2(d\theta^2 + \sin^2\theta d\phi^2) \\ &- 4\frac{m_{\rm BH}r^3\sqrt{\mathcal{K}}}{d^4}\left(1+\frac{\mathcal{K}}{d^2}\right)^{1/2} dt[\cos\theta\sin(\phi-\Omega t)d\theta + \sin\theta\cos(2\theta)\cos(\phi-\Omega t)d\phi] \\ &+ \frac{m_{\rm BH}r^2}{d^3}\left[\left(\frac{1}{2}+\frac{3\mathcal{K}}{2d^2}\right)(1-3\cos^2\theta) + \left(\frac{3}{2}+\frac{3\mathcal{K}}{2d^2}\right)\sin^2\theta\cos(2(\phi-\Omega t))\right] \\ &[dt^2 + dr^2 + r^2(d\theta^2+\sin^2\theta d\phi^2)]\,. \end{aligned} \tag{4.95}$$

4.4.2 Matching procedure

Our goal is to find a complete metric for the neutron star and its exterior which asymptotes to (4.95) at large r (i.e. $R \ll r \ll \mathcal{L}$). Looking at (4.95) carefully, we see that the first set of terms represents flat space, which is the asymptotic form of $\hat{g}_{\mu\nu}$, the Schwarzschild metric. The other terms are the perturbations. It is easy to see that $h_{t\theta}$ and $h_{t\phi}$, the products of magnetic tidal terms, have angular dependences corresponding to the odd-parity vector spherical harmonics $\Phi_{2(\pm1)j}$. (Note that these magnetic terms are purely relativistic.) The other perturbations, the products of elec-

tric tidal terms, have angular dependences corresponding to the scalar harmonics Y_{20}, $Y_{2(\pm2)}$ (for g_{tt}, g_{rr}) and the even-parity tensor harmonics Φ_{20jk}, $\Phi_{2(\pm2)jk}$ (for $g_{\theta\theta}, g_{\phi\phi}$). Therefore, the perturbations we are looking for are the $\ell = 2$ perturbations to a spherically symmetric spacetime. Furthermore, thanks to our choice of THZ coordinates [68], the only terms which appear are the terms allowed by the Regge-Wheeler gauge. We have already solved for the radial dependence of these perturbations; (4.95) just provides the final boundary condition.

It should be noted that this metric still does not take into account Alvi's *ad hoc* corrections for the finite size of the neutron star. It turns out that these corrections are irrelevant to us. The first adjusts the form of the magnetic, odd-parity terms $g_{t\theta}$ and $g_{t\phi}$, arguing that at lowest order in $m_{\rm BH}/d$, these terms should be a Lorentz boost of the electric, even-parity terms. This requires a modification of the boost velocity to account for the neutron star mass ($m_{\rm BH} \to m_{\rm BH} + m_{\rm NS}$). However, since we are not interested in the odd-parity perturbations in this work, these corrections are irrelevant.

The other correction Alvi makes is to Ω. When he later matches (4.95) to the post-Newtonian metric in region III, he finds a more accurate formula for the rotation rate. In our analysis, however, we will ignore the time dependence in (4.95). We are looking for *static* perturbations; the ultimate goal is to find the point of tidal disruption when no static solution can be found. Without loss of generality, we can take $t = 0$.

The complete procedure is now straightforward. Outside the star, the perturbations $H_0(r)$ and $K(r)$ are described by the analytic solutions (4.67) and (4.68), which are parameterized by the unknown constants c_1 and c_2. At large r, the first term of these expressions $\sim r^2$; matching to (4.95) gives

$$c_1^0 = -4\sqrt{\frac{\pi}{5}}\frac{m_{\rm BH}m_{\rm NS}^2}{d^3}\left(\frac{1}{2} + \frac{3\mathcal{K}}{2d^2}\right) \tag{4.96}$$

for the perturbations with $m = 0$ and

$$c_1^{\pm2} = 2\sqrt{\frac{2\pi}{15}}\frac{m_{\rm BH}m_{\rm NS}^2}{d^3}\left(\frac{3}{2} + \frac{3\mathcal{K}}{2d^2}\right) \tag{4.97}$$

for the perturbations with $m = \pm 2$. (As promised, we have restored the proper normalizations of the $Y_{\ell m}$.) The second term dies away at large r, so c_2 cannot be determined by this matching procedure. Note that in Alvi's case, the second term is poorly behaved at the horizon $r = 2m_{\rm NS}$, forcing $c_2 = 0$. This is not true in our case, since we have a neutron star for which $R > 2m_{\rm NS}$.[8]

Inside the star, equations (4.54) and (4.55) can be integrated to find $H_0(r)/c_0$ and $K(r)/c_0$, where c_0 is the unknown initial scaling of (4.60) and (4.61). The junction conditions tell us that both H_0 and K are continuous across the surface of the star. We can therefore equate the integrated $H_0(R)$ and $K(R)$ to the analytic solutions (4.67) and (4.68) evaluated at the same point. These two equations are sufficient to determine c_0 and c_2 and provide a complete solution for the metric.

4.5 Results: Tidal disruption

4.5.1 Procedural issues

The code we use computes the structure of a tidally perturbed star for two different choices of equation of state. The first is a standard polytrope,

$$P = K_p \rho_0^\gamma \,, \tag{4.98}$$

where ρ_0 is the *rest mass* density. In a relativistic situation, this must be distinguished from the total energy density ρ, which appears in the structure equations:

$$\rho = \rho_0 + \frac{P}{\gamma - 1} \,. \tag{4.99}$$

The second choice is a table listing values of P and ρ. In this case, the code determines ρ from a given P by using cubic spline interpolation [200].

The code begins by calculating the structure of the unperturbed star. The Tolman-

[8] Another difference from Alvi's work is that he uses isotropic coordinates for the perturbed black hole, rather than the Schwarzschild coordinates we use for the perturbed neutron star. This turns out not to matter for the matching, since all differences in the perturbations die off at large r.

Oppenheimer-Volkoff equations are integrated using an adaptive-step Runge-Kutta routine (just like the routine used in Chapters 2 and 3). At $r = 0$, $m = 0$. $\Phi(0) \neq 0$, but since we do not know its value, we set it to zero initially and then rescale it later. $\hat{P}(0)$ is obtained in one of two ways: For the polytrope case, it is calculated from an initial rest mass density. For the general equation of state, it is given directly as an input and the mass or baryon density does not appear at all. Since the TOV equations are poorly behaved at $r = 0$, we also evaluate m, $\hat{P}$, and Φ at small r using the expansions (4.56), (4.57), and (4.59). This helps get the integration started.

Some tabulated equations of state have density discontinuities due to phase transitions. In these cases, the integration briefly stops at the discontinuity, not taking any steps into the new regime. The integration concludes at the stellar surface $r = R$, where $\hat{P}(R) = 0$. In order to better handle these boundaries, which are defined by a certain value of pressure, we actually integrate an alternate form of the TOV equations in which pressure, not radius, is the independent variable. (That is, we integrate $dr/d\hat{P} = 1/(d\hat{P}/dr)$, $dm/d\hat{P} = (dm/dr)/(d\hat{P}/dr)$, and $d\Phi/d\hat{P} = (d\Phi/dr)/(d\hat{P}/dr)$.) Once the integration is complete, we can rescale $\Phi(r)$ by requiring that $\Phi(R) = \ln(1 - 2m_{\rm NS}/R)/2$, the Schwarzschild solution.

Next, the perturbed equations are integrated using the same adaptive-step Runge-Kutta routine. Since we now know the radii associated with density discontinuities and the surface, we can use radius as the independent variable. To start the integration, we use the expansions (4.60) and (4.61). When the integration is complete, we scale the perturbations to match the external solution, as described in the previous section.

Once the metric is obtained, the final step is to calculate the length of the principal axes of the star. To do so, we just integrate $\sqrt{g_{rr}}$ from the origin to the surface along a prescribed direction. However, as we noted earlier, the location of the surface is a gauge-dependent quantity. It makes sense to compute the axes in surface gauge. To do so, we need to adjust our coordinates in the radial direction by the *negative* of (4.76):

$$-\xi^r = \frac{\delta P}{\partial_r \hat{P}} = -\frac{H_{0m}}{2\Phi'} Y_{2m}(\theta, \phi)\,, \tag{4.100}$$

where H_{0m} is the perturbation associated with a particular m. (Remember we have $m = -2, 0, 2$.) Under a gauge transformation, the metric perturbation transforms like

$$h_{\mu\nu} \to h_{\mu\nu} - \hat{\nabla}_\mu \xi_\nu - \hat{\nabla}_\nu \xi_\mu \,, \tag{4.101}$$

so g_{rr} becomes

$$g_{rr} = e^{2\Lambda} \left[1 + \sum_{m=0,\pm 2} \left(H_{0m} + \frac{H_{0m}\Lambda'}{\Phi'} + \frac{H'_{0m}}{\Phi'} - \frac{H_{0m}}{\Phi'^2}\Phi'' \right) Y_{2m}(\theta, \phi) \right] , \tag{4.102}$$

The square root of this expression is integrated from 0 to R in three distinct directions: (1) $\theta = \pi/2$, $\phi = 0$ or π, the stretched axis pointing toward the black hole[9]; (2) $\theta = \pi/2$, $\phi = \pi/2$ or $3\pi/2$, the squeezed axis in the direction of the orbital motion; and (3) $\theta = 0$, the squeezed axis perpendicular to the orbit. We call these axes a_1, a_2, and a_3. It turns out that $a_1 > a_2 > a_3$. If we consider only the Newtonian tidal terms ($\mathcal{K} = 0$), $a_1 > a_2 = a_3$. When the perturbations become large enough, a_3 and then a_2 will cease to exist; functionally, this is because the expression under the square root becomes negative. The lack of a solution means that the star is tidally disrupted.

We consider *sequences* in which a single star is brought closer and closer to a black hole. Far away from the black hole, we choose a central density/pressure for the star and then integrate to find its structure. As we move closer and tidal effects become stronger, the central density/pressure may change. (In fact, it has been shown that the central density will decrease, stabilizing the star against collapse [151, 93].) We must therefore be careful that the star we are integrating at each orbital radius is the *same star*. The baryon mass, or rest mass, should remain constant:

$$m_b = \int \rho_0 dV_{\text{proper}} = \int \rho_0 \sqrt{g_{rr}} dr \sqrt{g_{\theta\theta}} d\theta \sqrt{g_{\phi\phi}} d\phi \,. \tag{4.103}$$

As we move closer to the hole, we can calculate the perturbed baryon mass as we go, adjusting the central density/pressure so that m_b remains constant. There are two sources of perturbations to m_b, in the density ρ_0 and in the metric functions. Both

[9]Recall we have chosen $t = 0$.

of these sources have the same angular dependence, and it can be shown that the perturbations average to zero when integrated over the sphere. The central density will *not* change in our formalism; indeed, Flanagan [93] has shown that this is a higher-order tidal effect. Therefore, we do not even need to calculate baryon mass. Instead we can parameterize a star by its unperturbed gravitational mass $m_{\rm NS}$, or better yet, its compactness $C = m_{\rm NS}/R$. (Note that R is the coordinate radius in Schwarzschild coordinates, not the physical lengths a_1, a_2, and a_3. This definition of compactness is consistent with the sources which we use for comparison.)

4.5.2 Comparison to other relativistic tidal problems

We would like to compare our results to those of Wiggins and Lai [254], who used a relativistic tide but Newtonian self-gravity. According to Wiggins and Lai's results, the relativistic tide causes earlier tidal breakup because "GR means 'stronger' gravity." They estimate that relativistic self-gravity will have a similarly sized effect, but in the opposite direction; in this case, the stronger gravity of GR helps keep the star together longer! For the comparison, we model the star as a polytrope with $K_p = 7.5 \times 10^4$ g^{-1} cm^5 s^{-2} and $\gamma = 2$ ($n = 1$). This is a reasonable approximation to the neutron star high-density equation of state.[10]

To perform the comparison, we actually use the results of Ferrari, Gualtieri, and Pannarale (FGP) [88]. This recent paper also tackles the problem of extending Wiggins and Lai to relativistic self-gravity; however, their method differs from ours. Their formalism is essentially the compressible Roche-Riemann model of Wiggins and Lai, but instead of using the Newtonian gravitational potential in these equations, FGP use a quasi-relativistic potential generated by the TOV equations for the unperturbed star. Our results, on the other hand, are completely relativistic, but the tides are represented by linear perturbations which may not be accurate in the strong-field regime. Table B1 of [88] compares the quasi-relativistic results to the results of Wiggins and

[10]This is a fit to Fig. 4 of [14], made by Jake Hartman for MIT subject 8.901. The equations of state we use later in this chapter come from a much newer paper [161] and seem to be a bit steeper; however, the authors of [161] claim that $n = 1$ is still a decent fit. In any case, a polytrope is not a great approximation.

C	q	Full GR	Quasi GR	Newtonian	Perturbative GR
0.1088	9	5.16	5.16	6.01	4.30
	8	5.56	5.48	6.47	4.50
	7.5	5.78	5.69	6.70	4.62
	7	6.02	5.89	6.91	4.76
	6	6.62	6.48	7.64	5.09
0.1201	7	5.49	5.42	6.52	4.44
	6.5	5.71	5.67	6.81	4.58
	6	6.01	5.93	7.15	4.73
	5	6.67	6.63	7.94	5.13
	3	9.47	9.03	10.9	6.58
0.1321	6	5.52	5.48	6.71	4.43
	5	6.21	6.10	7.46	4.77
	4	7.18	6.93	8.57	5.27
	3	8.62	8.25	10.29	6.06
0.1452	5	5.52	5.61	7.04	4.46
	4	6.46	6.39	8.09	4.89
	3	7.75	7.55	9.65	5.58
	2	10.4	9.71	12.5	6.86
	1	17.6	15.1	19.6	10.1
0.1600	4.5	5.52	5.50	7.15	4.34
	4	5.93	5.87	7.67	4.55
	3	7.19	6.92	9.16	5.15
0.1780	3.5	5.79	5.77	7.93	4.44
	3	6.38	6.30	8.71	4.73
	2	8.41	8.01	11.3	5.69

Table 4.2: Neutron star tidal disruption radius $r_{\rm tide}$, in units of black hole mass $m_{\rm BH}$, for various black hole-neutron star binaries and calculation methods. The binaries are parameterized by the neutron star's compactness $C = m_{\rm NS}/R$ and by the mass ratio $q = m_{\rm BH}/m_{\rm NS}$. "Full GR" results are calculated using numerical relativity [232]. "Quasi GR" results are calculated using a potential derived from the TOV equations [88]. "Newtonian" results use Newtonian self-gravity for the neutron star [254]. Finally, "perturbative GR" results are calculated using the methods of this chapter.

Lai (recalculated by FGP) as well as exact results computed in numerical relativity [232]. We reproduce that table as Table 4.2, including our own results for comparison. In doing so, we convert the FGP breakup frequency $\Omega_{\rm tide}$ to a breakup radius, $r_{\rm tide}$. FGP also use a system of units in which all quantities are scaled by $K_p^{n/2}$, whereas our $r_{\rm tide}$ is scaled by black hole mass $m_{\rm BH}$. The result is that the first three columns of Table 4.2 are found using

$$r_{\rm tide} = \left(\frac{1+q}{q^3 m_{\rm NS}^2 \Omega_{\rm tide}^2}\right)^{1/3} m_{\rm BH}\,, \tag{4.104}$$

where $q = m_{\rm BH}/m_{\rm NS}$.

Independent of method, we see that for a given compactness, tidal disruption happens earlier (larger $r_{\rm tide}$) for a smaller black hole. Smaller black holes have stronger tidal fields. In addition, tidal disruption happens earlier for less compact (smaller C) neutron stars; in these stars, the relativistic effects are not as strong as in stars with larger C.

Looking at the different methods, we see that the Newtonian method of Wiggins and Lai generally overestimates the tidal disruption radius by a significant amount; this is the effect of ignoring relativistic corrections to the neutron star's self-gravity. The quasi-relativistic potential method of FGP does much better. In certain situations (large q), it matches the numerical results very well. When it deviates from the numerical results, it does so in the opposite way from the Newtonian method: The disruption radius is too small.

Our own results also tend to underestimate the disruption radius. Again, the agreement is better for larger q and worse for smaller q. However, our underestimate is much worse than that of FGP. In most cases, our results for $r_{\rm tide}$ are $\sim 70-80\%$ smaller than the numerical results. In extreme cases, this can get as low as $\sim 60\%$. It seems, therefore, that our method for estimating tidal disruption is really no better than the Newtonian self-gravity method; it just errs in the opposite direction. While we incorporate the relativistic self-gravity of the neutron star, our approximations are clearly too strong. One assumption that jumps out is the use of circular orbits

to find the tidal field, when many of our results extend past the ISCO. Still, we should expect to match the FGP results better in these cases, since they use the same method to calculate the tide. Perhaps the biggest approximation is the basic assumption that perturbation theory—technically valid only for small perturbations to the metric—can be used to describe a strongly distorted, disrupting star.

4.5.3 Comparison to classical tidal problems

A useful way to check the validity of the perturbative method is to consider a white dwarf instead of a neutron star. In this case, the gravity of the star is almost Newtonian, so we can consider compare our results to those found by Chandrasekhar for the classical Jeans and Roche problems. To do so, we must also set $\mathcal{K} = 0$ so that the black hole produces only a Newtonian tidal field. With this approximation, the tidal field is the same in the y and z directions, so $a_2 = a_3$. For a white dwarf binary, the orbit is not relativistic enough for the term with $\mathcal{K}$ to matter much anyway.

For a white dwarf, the equation of state is well-described by a polytrope of $\gamma = 5/3$ ($n = 1.5$). We choose $K_p = 3.1232 \times 10^{12}$ g$^{-2/3}$ cm^4 s^{-2} and central density $\rho_{0c} = 10^6$ g cm^{-3}. With these choices, we find $m_{\rm NS} = 0.49 M_\odot$ and $R = 1.1 \times 10^4$ km.

For $n = 0$, we could have just read off the exact solutions from Chandrasekhar's book (Tables XVI and XVII for the Roche and Jeans problems, respectively). Each row gives axis ratios a_2/a_1 and a_3/a_1 (for Roche) or eccentricity e (for Jeans) and the associated value of $x_{\rm Chandra} \equiv \psi/\pi\rho$ that corresponds to that deformation. (Table XVI actually gives $(1 + 1/q)x_{\rm Chandra}$.) Given the $m_{\rm NS}$ and R of our star and a particular $m_{\rm BH}$ for the black hole, we could convert $x_{\rm Chandra}$ to an orbital radius and then plug that radius into our code to see what sort of deformation it gives at the same part of the inspiral.

Since $n > 0$, we have to use the compressible form of Chandrasekhar's equations derived by Lai, Rasio, and Shapiro [152, 153] and used in Wiggins and Lai [254]. It turns out that these equations are identical to (4.1) and (4.2) (Roche) or (4.14) (Jeans) with two simple modifications. First, (4.1), (4.2), and (4.14) contain the density ρ, which for an $n = 0$ polytrope is constant: $\rho = 3m_{\rm NS}/4\pi R^3$. In the compressible case,

this is replaced by the average density $\bar{\rho} = 3m_{\rm NS}/(4\pi a_1 a_2 a_3)$. Second, the left-hand sides of these equations must be multiplied by $\kappa_n(1 - n/5)$, where κ_n is a constant that characterizes the star's moment of inertia [152]. For $n = 1.5$, $\kappa_n = 0.51149$.

These modifications to the Roche and Jeans equations amount to the transformation $\psi/\pi\rho \to \psi\kappa_n(1 - n/5)/\pi\bar{\rho}$. With this knowledge, we can make use of Chandrasekhar's tabulated results after all. We simply follow the same procedure described above, but we treat the quantities in the tables as $x_{\rm Chandra} = \psi\kappa_n(1 - n/5)/\pi\bar{\rho}$. However, things are more difficult for a compressible star: $\bar{\rho}$ depends on a_1, and the tables only give the axis ratios a_2/a_1 and a_3/a_1 (or in the Jeans case, e). In the compressible case, we actually need *three* equations to solve the Roche problem completely, while two are needed for the Jeans problem. (In the incompressible case, one of these equations was redundant.) This additional equation is:

$$-2\pi a_3 A_3 \bar{\rho} + \frac{m_{\rm NS}}{R a_3}\left(\frac{R}{(a_1 a_2 a_3)^{1/3}}\right)^{3/n} - \psi\kappa_n\left(1 - \frac{n}{5}\right)a_3 = 0\,. \tag{4.105}$$

This is the equation of motion for a_3 (with $\ddot{a}_3 = 0$); it corresponds to Wiggins and Lai equation (33). We could have chosen any of the three equations of motion given in that paper; however, a_3 is a bit simpler than the others because it is independent of the mass ratio parameter q. Solving this equation for a_1 and then substituting back into the expression for $\bar{\rho}$, we find for $n = 1.5$

$$\bar{\rho} = \rho\frac{\hat{a}_3^4}{\hat{a}_2^2}\left[\frac{3}{2}\left(A_3 + \frac{x_{\rm Chandra}}{2}\right)\right]^3\,, \tag{4.106}$$

where $\hat{a}_2 = a_2/a_1$ and $\hat{a}_3 = a_3/a_1$. Given $\bar{\rho}$, we can now convert $x_{\rm Chandra}$ to an orbital radius.

Figure 4-1 shows the comparison between the perturbative GR results and those calculated in the Jeans problem. For this case, we use a black hole mass $M = 3\ M_\odot$. At large r, the two methods agree quite well. As r decreases, the perturbative method tends to overestimate the tidal squeezing. (It should be noted, however, that in the compressible case, the Jeans solution is itself not exact.) Eventually, the two solutions

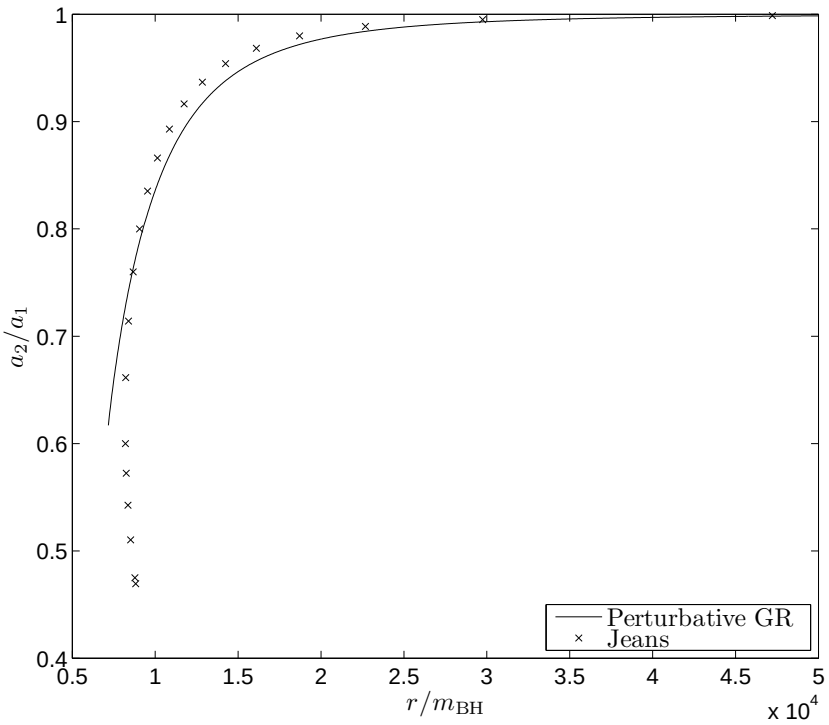

Figure 4-1: Comparison of perturbative GR results and the classical Jeans results (modified for a compressible star) for the tidal distortion and disruption of a 0.49 $M_\odot$ white dwarf orbiting a 3 $M_\odot$ black hole.

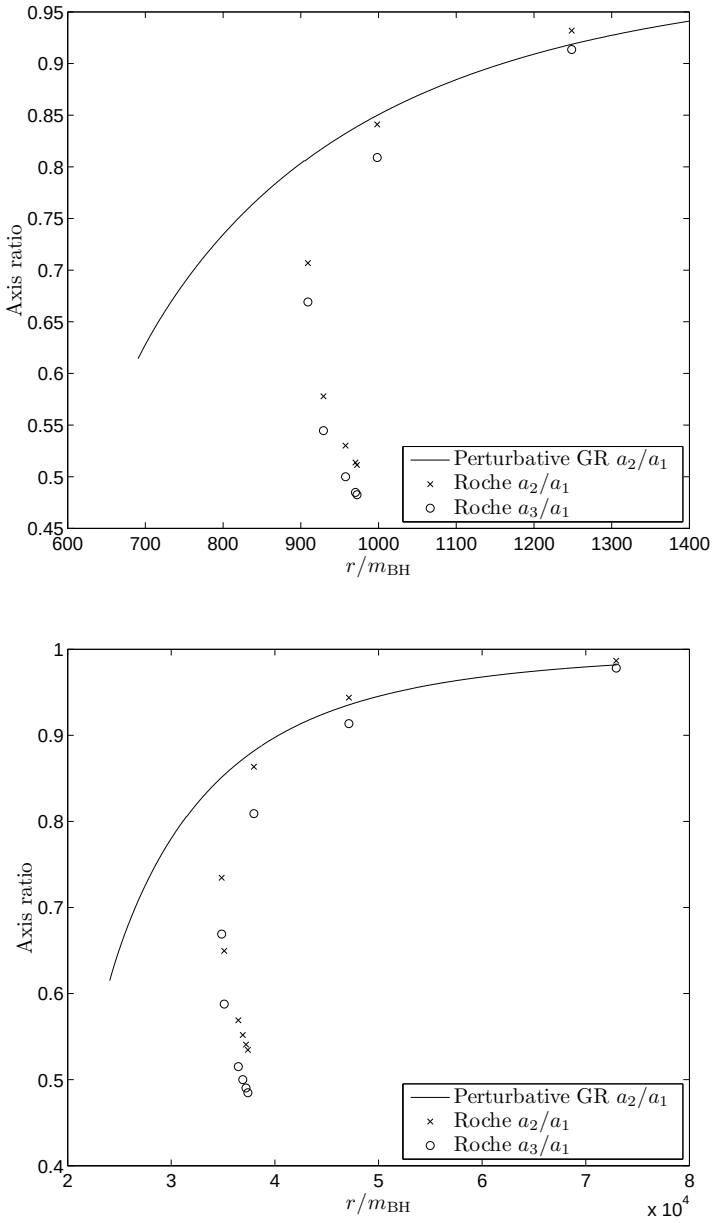

Figure 4-2: Comparison of perturbative GR results and the classical Roche results (modified for a compressible star) for the tidal distortion and disruption of a 0.49 $M_\odot$ white dwarf orbiting a black hole. The top plot shows the case $m_{\rm BH} = 100\ M_\odot$ ($q \to \infty$), while the bottom plot shows the case $m_{\rm BH} = 0.49\ M_\odot$ ($q = 1$)

meet again, and the perturbative method begins to underestimate the tidal effect.

The Jeans solution is plotted using only the points which lie below the $n = 0$ Roche limit. This limit occurs when $x_{\rm Chandra}$ is maximum; for $n = 0$, that corresponds to the minimum possible orbital radius. In the compressible case, $x_{\rm Chandra}$ also contains a dependence on $\bar{\rho}$, so its maximum does not represent the minimum radius. This causes the Jeans curve to turn around. We define the Roche limit to occur instead when $r_{\rm orb}$ is minimal. (For our own results, breakup occurs when one of the axis ratio integrands fails to exist; therefore, there can be no points beyond breakup.) We see that the perturbative breakup happens at a slightly smaller radius than the Jeans breakup (approximately 87% smaller, though the Jeans points are too separated to be accurate).

Even though the two curves and their respective tidal breakup points do not match, the agreement is still relatively impressive considering that perturbation theory is technically not appropriate to use for strong tides. The agreement is not so good, however, when we compare to the Roche problem. Figure 4-2 shows these results for the two choices $m_{\rm BH} = 100\ M_\odot$ (large enough to compare to the $q = \infty$ Roche results) and $m_{\rm BH} = m_{\rm NS}$ ($q = 1$). In both cases, the perturbative distortion agrees with the classical result at large radius, but the two results diverge strongly at smaller radius. The perturbative breakup radius is $\approx 76\%$ (69%) smaller for the $q \to \infty$ ($q = 1$) case. This level of disagreement is on par with the underestimates seen in Table 4.2.

Since we do not include the rotation of the star in our model, the Jeans comparison is the best test of how well the perturbative method works for strong tides: not well, but not as bad as might be expected. The Roche comparison, however, tells us that ignoring rotation, and the additional distortion it causes, is fatal when trying to compare to results which include this effect. In fact, the other results in Table 4.2 do not even correspond to the Roche problem. Instead, they are extensions of the Roche-Riemann problem, which also includes internal fluid motions. Ignoring these effects must also contribute to some of the disagreement in Table 4.2.

In the end, it seems that perturbation theory, at least with the approximations

and limitations we have introduced, is probably not the best way to study strong tidal effects such as disruption. Although it is possible this method could be improved, the rest of this chapter will instead change gears and focus on a different application of this formalism, the calculation of relativistic Love numbers.

4.6 Love numbers

Consider a static, spherically symmetric star placed in a time-independent quadrupolar tidal field $\mathcal{E}_{ij}$. The tide will induce a quadrupole moment Q_{ij} in the star.[11] To lowest order, the quadrupole moment will be proportional to the tidal field [95]:

$$Q_{ij} = -\lambda \mathcal{E}_{ij} \,. \tag{4.107}$$

The (dimensionless) Love number is then defined as

$$k_2 \equiv \frac{3}{2} \lambda R^{-5} \,, \tag{4.108}$$

where R is the radius of the star. The Love number is also often known as the apsidal constant. It characterizes how a star responds to tidal perturbations in general, without reference to a specific tidal potential.

As a neutron star orbits another neutron star, the companion's tidal field excites its normal modes of oscillation at the orbital frequency Ω, much below the modes' resonant frequencies. This excitation alters the total quadrupole moment and energy of the binary system, leading to corrections to the gravitational wave phase. To lowest order, these corrections can be directly related to the Love number and no other parameters (besides the masses, radii, and frequency). Even though this correction, by construction, is only valid for small tides during the early phases of inspiral, it may have an observable effect. Tidal effects become more complicated near merger; many previous investigations (e.g., [82]) have studied this regime.

[11] In Chapter 1, we called this $\mathcal{I}_{ij}$, but we change notation here to match Flanagan and Hinderer [95].

In this section, we first derive the tidal impact on GWs, following the work of Flanagan and Hinderer [95] very closely. Their result initially depends also on the mode frequencies ω_n, but in the end they make the approximation $\Omega \ll \omega_n$ and recover the results of [178], which depend only on the Love number. Then we discuss how to compute Love numbers for a star with relativistic self gravity. It turns out that the Love number can be easily obtained using our previously developed tidal disruption code.

4.6.1 Effect of tidal excitations on GW phase

To obtain the tidal influence on the gravitational wave phase, Flanagan and Hinderer use a quasi-Newtonian formalism. Following them, we first write down a Lagrangian for the system. For a star of mass m_1 in a circular orbit around another star of mass m_2, we find

$$L = \frac{1}{2}\mu\dot{r}^2 + \frac{1}{2}\mu r^2\dot{\Phi}_{\rm orb}^2 + \frac{M\mu}{r} - \frac{1}{2}Q_{ij}\mathcal{E}_{ij} + \sum_n(\alpha_n\dot{Q}^n_{ij}\dot{Q}^n_{ij} - \beta_n Q^n_{ij}Q^n_{ij})\,, \qquad (4.109)$$

where M is the total mass, μ is the reduced mass, $\Phi_{\rm orb}$ is the phase of the orbit, and α_n and β_n are currently unknown constants. The first three terms are the standard Lagrangian of an orbiting point mass. The fourth term is an interaction between the induced quadrupole and the tidal field. The final two terms are due to the oscillation of modes with n radial nodes. The form of these terms can be derived as follows [94]: In general, there are five modes with $\ell = 2$ and n radial nodes. We can write the Lagrangian in terms of the mode excitations, with kinetic and potential pieces. However, we can also map the excitations to the five independent pieces of the induced quadrupole Q^n_{ij} and then write the Lagrangian in terms of that quantity and its time derivative. Due to rotational invariance, only traces of products of these matrices can appear in the Lagrangian. This sets the general form of (4.109).

We now determine the constants α_n and β_n [94]. The ratio of the two constants $\beta_n/\alpha_n = \omega_n^2$, so the problem really reduces to determining the overall scaling. To do this, imagine putting the star in a static tidal field. Then vary the action with respect

to Q^n_{ij} to get the Euler-Lagrange equation

$$2\alpha_n \ddot{Q}^n_{ij} + 2\omega_n^2 \alpha_n Q^n_{ij} + \frac{1}{2}\mathcal{E}_{ij} = 0\,. \tag{4.110}$$

In a static field, the first term is zero. Solving for α_n using (4.107), we find $\alpha_n = (4\lambda_{1,n}\omega_n^2)^{-1}$. $\lambda_{1,n}$ is the contribution to the Love number of star 1 if only the modes with n radial modes are excited. The total Love number $\lambda_1 = \sum_n \lambda_{1,n}$.

Now that the full Lagrangian has been determined, we return to the case of an orbiting star. The (Newtonian) tidal field can be written as $\mathcal{E}_{ij} = -m_2\partial_i\partial_j(1/r)$. Then the complete Euler-Lagrange equations can be written

$$\ddot{x}^i + \frac{Mx^i}{r^3} = \frac{m_2}{2\mu} Q_{jk}\partial_i\partial_j\partial_k \frac{1}{r}\,, \tag{4.111}$$

$$\ddot{Q}^n_{ij} + \omega_n^2 Q^n_{ij} = m_2\lambda_{1,n}\omega_n^2\partial_i\partial_j\frac{1}{r}\,, \tag{4.112}$$

where $Q_{ij} = \sum_n Q^n_{ij}$. These equations do not include radiation reaction. These terms could be included and the equations numerically integrated. However, Flanagan and Hinderer choose instead to find equilibrium circular orbit solutions with $r = $ const. and $\Phi_{\rm orb} = \Phi_{\rm orb,0} + \Omega t$. Linearizing in $\lambda_{1,n}$, we can find the radius as a function of frequency:

$$r = M^{1/3}\Omega^{-2/3}\left[1 + \frac{3}{4}\sum_n \chi_n g_1(x_n)\right]\,, \tag{4.113}$$

where $\chi_n = m_2 m_1^{-1} M^{-5/3}\Omega^{10/3}\lambda_{1,n}$, $g_1(x) = 1+3(1-4x^2)^{-1}$, and $x_n = \Omega/\omega_n$. Here the first term is the standard Kepler result, and the second term is the linear correction due to the tidal excitation. By switching some signs in the Lagrangian, we can also find the energy of the system:

$$E = -\frac{\mu}{2}(M\Omega)^{2/3}\left[1 - \frac{9}{4}\sum_n \chi_n g_2(x_n)\right]\,, \tag{4.114}$$

where $g_2(x) = 1 + (3-4x^2)(1-4x^2)^{-2}$. To find the energy loss due to gravitational waves, we first define the *total* quadrupole moment Q^T_{ij} as the sum of Q_{ij} and the

orbital component: $Q_{ij}^T = Q_{ij} + \mu x_i x_j - \mu r^2 \delta_{ij}/3$. Then the energy loss is given by the quadrupole formula:

$$\frac{dE}{dt} = -\frac{1}{5}\langle \dddot{Q}_{ij}^T \dddot{Q}_{ij}^T \rangle = -\frac{32}{5} M^{4/3} \mu^2 \Omega^{10/3} \left[1 + 6 \sum_n \chi_n g_3(x_n) \right] , \tag{4.115}$$

with $g_3(x) = (M/m_2 + 2 - 2x^2)(1 - 4x^2)^{-1}$. Again, the first term is the standard result; we are interested in the second term, which contains the tidal correction.

As in Chapter 2, we can combine the energy E with the energy loss dE/dt to find the frequency chirp $df/dt = (dE/dt)/(dE/df)$ (where $f = \Omega/\pi$ is, as before, the gravitational wave frequency). The gravitational wave phase is then $\Phi = \int 2\pi f dt$. In a complete analysis of the GW tidal effect, we would calculate Φ (or, as in Chapter 2, its stationary phase counterpart Ψ) and then compute the Fisher matrix, including the Love number as an additional parameter. In order to just get a quick idea of the term's impact, we can simply compute the additional number of cycles it contributes to the GW phase. This is the accumulated phase divided by 2π:

$$N_{\rm cyc} = \int_{f_{\rm low}}^{f_{\rm high}} f dt = \int_{f_{\rm low}}^{f_{\rm high}} f \frac{dE}{df} \left(\frac{dE}{dt} \right)^{-1} df \, . \tag{4.116}$$

If we evaluate this at lowest order in λ, we find

$$N_{\rm cyc} = \int_{f_{\rm low}}^{f_{\rm high}} \frac{5}{96 \mu M^{2/3} \Omega^{8/3}} \left[1 - \frac{1}{3} \sum_n \chi_n g_4(x_n) \right] df \, , \tag{4.117}$$

where

$$g_4(x) = \frac{2M}{m_2(1-4x^2)} + \frac{22 - 117x^2 + 348x^4 - 352x^6}{(1-4x^2)^3} \, . \tag{4.118}$$

We now finally simplify by assuming $x_n \ll 1$. The final integrated result for the *change* in cycles due to the tidal term is

$$\Delta N_{\rm cyc} = -\frac{3\pi^{2/3}}{16 \mu M^{7/3}} \left[\left(11 \frac{m_2}{m_1} + \frac{M}{m_1} \right) \lambda_1 + \left(11 \frac{m_1}{m_2} + \frac{M}{m_2} \right) \lambda_2 \right] (f_{\rm high}^{5/3} - f_{\rm low})^{5/3} \, , \tag{4.119}$$

including the contribution from the second star ($1 \leftrightarrow 2$) as well. If we consider an

equal mass binary, in which both stars have mass $m_{\rm NS}$, the result is

$$\Delta N_{\rm cyc} = -\frac{39}{16}\left(\frac{\pi^2}{2}\right)^{1/3}\frac{\lambda}{m_{\rm NS}^{10/3}}(f_{\rm high}^{5/3} - f_{\rm low})^{5/3}\,. \tag{4.120}$$

4.6.2 Calculating the Love number

For Newtonian stars, the Love number can be calculated using solutions of the Clairaut-Radau differential equation [40]. However, as we now know quite well, neutron stars require a *relativistic* treatment of tidal effects. Berti, Iyer, and Will [31] recently approximated the relativistic Love number by calculating the quadrupole induced by *rotation* in the slow-rotation framework of Hartle and Thorne [107, 110]. In the Newtonian case, the $\ell = 2$ rotational apsidal constant is equivalent to the $\ell = 2$ tidal apsidal constant. In GR, the cases are somewhat different due to frame-dragging effects and the fact that rotational energy gravitates. However, Berti, Iyer, and Will believe that the rotational result should be good enough to approximate the Love number. Hinderer, on the other hand, has calculated the relativistic Love numbers directly using $\ell = 2$ tidal perturbations [116]. We describe that procedure here.

In the LARF of a neutron star, the tt piece of the metric can be expanded as [238]

$$g_{00} = -1 + \frac{2m_{\rm NS}}{r} + 3\frac{Q_{ij}n^in^j}{r^3} - \mathcal{E}_{ij}n^in^jr^2\,, \tag{4.121}$$

where $m_{\rm NS}$ is the mass of the star, Q_{ij} is its (symmetric, trace-free) quadrupole moment, and $\mathcal{E}_{ij}$ is the electric tidal field of the other neutron star. By (4.107), $Q_{ij} = -\lambda\mathcal{E}_{ij}$. If we change basis from symmetric-trace-free (STF) tensors [237] to spherical harmonics, g_{00} can instead be expanded as

$$g_{00} = -1 + \frac{2m_{\rm NS}}{r} + \frac{3Q_m}{r^3}Y_{2m}(\theta,\phi) - \mathcal{E}_m r^2 Y_{2m}(\theta,\phi)\,. \tag{4.122}$$

In this basis, we have $Q_m = -\lambda\mathcal{E}_m$. We can calculate the $\ell = 2$ static perturbations to a neutron star and match the form of the solution at large r to (4.122). Q_m,

$\mathcal{E}_m$, and λ are then easily read off.[12] In effect, this procedure is the same as that described in the first part of this chapter. However, earlier we were trying to match to a *specific* tidal field of a black hole. In this case, we are matching to a *general* field and calculating the intrinsic response of the star. Aside from this small difference, all of the necessary machinery has already been built.

Hinderer herself uses a second-order equation for H_0:

$$\begin{aligned} &H_0'' + H_0' \left\{ \frac{2}{r} + e^{2\Lambda} \left[\frac{2m}{r^2} + 4\pi r(\hat{P} - \hat{\rho}) \right] \right\} \\ &\quad + H_0 \left[-\frac{6e^{2\Lambda}}{r^2} + 4\pi e^{2\Lambda} \left(5\hat{\rho} + 9\hat{P} + \frac{\hat{\rho} + \hat{P}}{dP/d\rho} \right) - 4\Phi'^2 \right] = 0 \,. \end{aligned} \tag{4.123}$$

This equation can be derived from our own equations (4.54) and (4.55). Outside the star, it reduces to (4.64). (Note that Hinderer's small r expansion of this equation, her Eq. (16), is not equivalent to our own expansion (4.60). Our expansion is correct, but the discrepancy should make little difference.) At first, it may seem that (4.123) is more valuable than (4.54) and (4.55) for the calculation of Love numbers since it eliminates any reference to $K(r)$, unneeded in this computation. But (4.54) and (4.55) are actually more useful in case of a density discontinuity inside the star. In (4.123), the term involving $dP/d\rho$ diverges at any such discontinuity.

Taking the solution outside the star, (4.67), we find the large r behavior:

$$H_0(r) = A \left(\frac{r}{m_{\rm NS}} \right)^2 + O\left(\frac{r}{m_{\rm NS}} \right) + \frac{8}{5} B \left(\frac{m_{\rm NS}}{r} \right)^3 + O\left(\left(\frac{m_{\rm NS}}{r} \right)^4 \right) . \tag{4.124}$$

Matching to (4.122), we find

$$A = -m_{\rm NS}^2 \mathcal{E}_m \,, \tag{4.125}$$

$$B = -\frac{15}{8 m_{\rm NS}^3} \lambda \mathcal{E}_m \tag{4.126}$$

[12] There may be some gauge ambiguity in this definition; it is worth investigating this in future work.

so that

$$\lambda = \frac{8}{15} m_{\rm NS}^5 \frac{B}{A} \tag{4.127}$$

or

$$k_2 = \frac{4}{5} \left(\frac{m_{\rm NS}}{R}\right)^5 \frac{B}{A} \,. \tag{4.128}$$

If we differentiate (4.67), we can relate B/A to H_0'/H_0 at some radius. The final result is

$$\begin{aligned} k_2 = {} & \frac{8C^5}{5}(1-2C)^2[2+2C(y-1)-y]\{2C[6-3y+3C(5y-8)] \\ & + 4C^3[13-11y+C(3y-2)+2C^2(1+y)] \\ & + 3(1-2C)^2[2-y+2C(y-1)]\ln(1-2C)\}^{-1} \,, \end{aligned} \tag{4.129}$$

where $C = m_{\rm NS}/R$ is the compactness of the star and $y = RH_0'(R_+)/H_0(R_+)$.

The complete procedure is now obvious. The unperturbed star gives us C. We then calculate H_0 starting at the center of the star, again with some arbitrary scaling. If necessary, we apply junction conditions at the surface to find the derivatives at $r = R_+$. This gives us y, and k_2 follows.

4.7 Results: Love numbers

4.7.1 Polytropes

Love numbers were first calculated for Newtonian polytropic models by Brooker and Olle [40]. A general result of this work is that Love numbers are largest for stiff polytropes (smaller n) and smallest for soft polytropes (larger n). This can be explained by the degree to which mass is centrally condensed in these polytropes. Stiff polytropes change their pressure more when density is changed, so the star resists becoming centrally condensed. The most extreme example is the $n = 0$ polytrope, which is a constant density star. On the other hand, in soft polytropes the pressure does not react as strongly to increases in density, so the mass sinks to the center easily. The high Love numbers in stiff stars can be explained by the fact that more

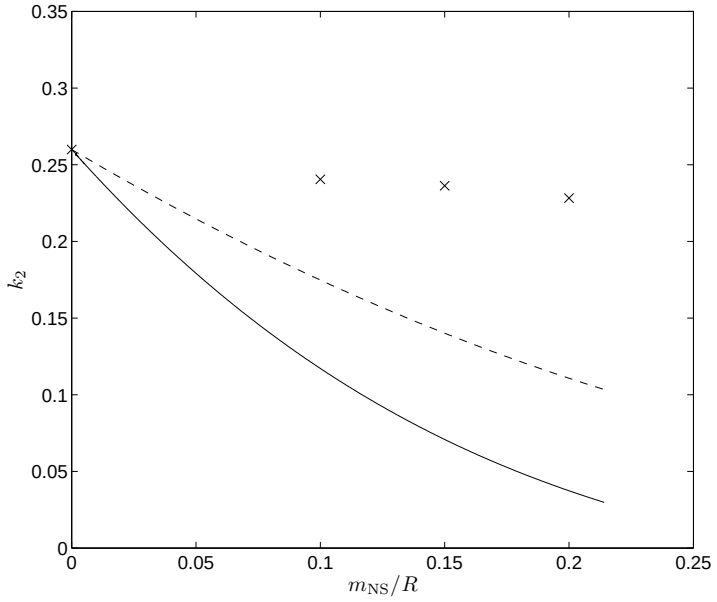

Figure 4-3: Love number k_2 as a function of compactness $m_{\rm NS}/R$. The solid line depicts k_2 for the equation of state $P = 7.4 \times 10^4 \rho_0^2$. The points marked $\times$ are Hinderer's original, incorrect results for a $\gamma = 2$ ($n = 1$) polytrope [116]. The dashed line is the $n = 1$ result from Berti, Iyer, and Will [31], who use a rotating star to approximate a tidally deformed one.

mass at large radii leads to a larger quadrupole moment; the reverse is true for a soft, centrally condensed star.

Hinderer [116] computed relativistic Love numbers for polytropes using the matching method described above. She found that the relativistic Love numbers did not differ too much from the Newtonian case; typically, they were smaller by $\sim 10 - 20\%$. In some (small n) cases, though, the relativistic Love numbers were actually a bit *larger* than in the Newtonian case. This result is counter-intuitive, since we expect a relativistic star to be less susceptible to tidal forces. In addition, we found a discrepancy with Hinderer's results while testing our tidal disruption code with an $n = 1$ polytrope. While we agreed quite well with Hinderer's (and Brooker and Olle's) Newtonian Love numbers, our relativistic Love numbers were much smaller. Eventually, we found the source of the discrepancy: Hinderer made several small, but critical, typographical errors while deriving her Eq. (23) (equivalent to our (4.129)).

Figure 4-3 shows the Love numbers calculated by Hinderer for a $\gamma = 2$ ($n = 1$) polytrope and the Love numbers calculated by our code for the same model. (For our results, $K_p = 7.5 \times 10^4$ g^{-1} cm^5 s^{-2} as before.) Note first that both studies agree on the value as $m_{\rm NS}/R \to 0$; this value equals that given for a Newtonian $n = 1$ polytrope in [40]. At larger values of the compactness $m_{\rm NS}/R$, we see that the true Love numbers are a factor of $2-5$ smaller than those calculated by Hinderer. Hinderer has since corrected her results [117]. They do remain slightly incorrect due to the use of total energy density ρ instead of rest mass density ρ_0 in the polytrope law, but for the most part, they now agree with our results. In addition, the relativistic Love number is now always smaller than the Newtonian case.

Figure 4-3 also shows the Love numbers derived by Berti, Iyer, and Will using a rotating star instead of a tidally deformed one [31]. We see that while their analysis captures the basic idea of the relativistic correction, the actual magnitude of the correction is underestimated.

The punchline is that, in keeping with the theme of this thesis, relativistic effects are quite important in determining neutron star Love numbers. The downside is that relativity makes the neutron stars *more resistant* to tidal deformation than originally thought. The already-poor ability to measure the Love number using gravitational waves, as calculated in [95] using the faulty values, must now be treated as overly optimistic.

4.7.2 General equations of state: definitions

As mentioned above, polytropes are a poor approximation to the complicated high-density neutron star equation of state. It is interesting to see how real equations of state change the conclusions of the (corrected) work by Hinderer. In this subsection, we describe the many equations of state we use; in the next, we present the final results.

The real equations of state can be broken into several regimes. The provided numerical data is often not very smooth across the transitions between regimes, causing problems for the cubic spline interpolation. Therefore, the code actually interpolates

within each regime separately from the others.

At low densities, we follow the strategy of [161] and, for most stars, use a common set of well-established equations of state. (The exception is for stars with quark matter at the surface, which are described in more detail later.) Below $n < 0.001$ fm^{-3}, we use tabulated results from Baym, Pethick, and Sutherland [23]. At the lowest densities, 7.86 g cm$^{-3} < \rho < 1150$ g cm^{-3}, these authors simply quote the previously calculated results of Feynman, Metropolis, and Teller (FMT) [89]. This equation of state uses the Thomas-Fermi-Dirac atomic model to represent electrons in a lattice of iron nuclei [221]. For pressures below the lowest tabulated FMT pressure, we always use the lowest tabulated FMT density. This gives a density jump at the surface, but it is small enough to have no discernible effect on the perturbations.

At higher densities, but below neutron drip, $\rho < 4.3 \times 10^{11}$ g cm^{-3}, the calculation in Baym, Pethick, and Sutherland is new to that paper. (Therefore, we shall call this EOS segment, and this segment only, BPS.) In this regime, inverse beta decay makes the nuclei more neutron rich [221]. As the density increases, different nuclei become stable. (At each transition, there is a slight density discontinuity, but these are washed out in the tabulated data.) Finally, above neutron drip, Baym, Pethick, and Sutherland quote the earlier work of Baym, Bethe, and Pethick (BBP) [22].

For 0.001 fm$^{-3} < n < 0.08$ fm^{-3}, we use the results of Negele and Vautherin (NV) [182], which is an improvement on the results of Baym, Bethe, and Pethick for the region beyond neutron drip but before nuclear density. The equation of state can be determined using the total energy per baryon, which in MeV is

$$E_T = m_n + c_0 + \exp\left(\sum_{i=1}^{7} c_i [\ln(10^{-35} n_b)]^{i-1}\right) . \tag{4.130}$$

m_n is the neutron mass in MeV, n_b is the baryon density in cm^{-3}, and the coefficients c_i can be found in [182]. From E_T we can obtain $P = n_b^2(\partial E_T/\partial n_b)$ and $\rho = n_b E_T/c^2$.

Baryon densities $n > 0.08$ fm^{-3} bring us to nuclear density and beyond. In this regime, the equation of state is most uncertain; it is also this regime that has the dominant effect on bulk properties of a neutron star. We use a selection of equations

of state generously provided in table format by Jim Lattimer and Mark Alford.[13] The equations provided by Lattimer are a subset of those used in [161]. Most of these contain only nucleons (neutrons and protons) and leptons:

- **AP1–AP4**: Taken from Akmal, Pandharipande, and Ravenhall [3, 4]. All four come from variational calculations using the Argonne v_{18} two-nucleon interactions; AP1 is the most basic example. AP2 includes relativistic boost corrections. AP3 includes the Urbana IX three-nucleon interactions, but not the relativistic corrections. AP4 includes all three effects and is the most accurate equation of state in the series.

- **ENG**: Taken from Engvik et al. [77]. The equation of state is calculated using a (relativistic) Dirac-Brueckner Hartree-Fock method and a Bonn meson-exchange potential.

- **FSU**: Taken from Todd-Rutel and Piekarewicz [243], this is a relativistic field model.

- **MPA1**: Taken from Müther, Prakash, and Ainsworth [180]. Similar to ENG, but older, this EOS is also based on a Dirac-Brueckner Hartree-Fock method with a Bonn potential.

- **MS1–MS3**: Taken from Müller and Serot [179], these are field theoretical calculations. MS1 corresponds to $\xi = \zeta = 0$, where ξ and ζ are various interaction strengths defined in [179]. MS2 corresponds to $\xi = 1.5$ and $\zeta = 0.06$. Finally, MS3 corresponds to MS1 with the symmetry energy changed to 25 MeV.[14]

- **PAL6**: Taken from Prakash, Ainsworth, and Lattimer [199]. This is one of a set of equations calculated using a phenomenological nonrelativistic potential.

[13] Thanks also to Josiah Schwab, who first obtained all the tabulated EOS described in this section and then passed them along for use in this project.

[14] Note that this information comes from the file header and does *not* seem to match up to the description of MS3 in [161]. However, the mass-radius relation matches MS3 in [163]. Note also that after the initial definitions of these cases in [161], the authors seem to change notation by subtracting one from each digit. For example, MS1 is called MS0 in several places throughout the paper; at each of these places, MS2 becomes MS1, and MS3 becomes MS2.

In this case, the bulk nuclear matter incompressibility $K_s = 120$ MeV. This makes it the softest EOS made up of pure nucleons.

- **WFF1–WFF3**: Taken from Wiringa, Fiks, and Fabrocini [258]. This is an earlier variational calculation. WFF1 uses the Argonne v_{14} two-nucleon potential with the Urbana VII three-nucleon potential. WFF2 and WFF3 both use the Urbana v_{14} two-nucleon potential, but WFF2 uses the Urbana VII three-nucleon potential, while WFF3 uses the TNI three-nucleon interaction of Lagaris and Pandharipande [150].

The following equations of state contain more exotic components:

- **PS**: Taken from Pandharipande and Smith [188], this equation of state only contains neutrons, not protons. It also contains a pion condensate.

- **GM1–GM3**: Taken from Table II of Glendenning and Moszkowski [101]. They use a field-theoretical model which includes hyperons in addition to nucleons.

- **GS1–GS2**: Taken from Glendenning and Schaffner-Bielich [102]. They use a field-theoretical model which includes kaons in addition to nucleons. GS1 corresponds to "GL78" in the original paper, with $U_K(\rho_0) = -140$ MeV. GS2 corresponds to "TM1" with $U_K = -185$ MeV.

- **PCL2**: Taken from Prakash, Cooke, and Lattimer [198], this is a field-theoretical model which also includes hyperons and free quarks. The quarks are represented with an MIT bag model in which the bag constant $B = 200$ MeV fm^{-3}.

Finally, the **SQM1–SQM3** equations of state consist *entirely* of quark matter made of up, down, and strange quarks ("strange quark matter"). It has been conjectured that strange quark matter is the ground state of matter [259]. If this is true, normal matter will spontaneously convert to strange quark matter at sufficiently high density. A SQM star is self-bound; it does not require gravity to hold itself together. Because the entire star is made of strange quark matter, we do not need to use the low-density equations of state in this case; the SQM equation of state works down to

zero pressure. In addition, the SQM EOS reaches zero pressure at a density above nuclear density, creating an extremely large discontinuity at the surface. SQM1 has bag constant $B = 94.349$ MeV fm^{-3}, strange mass $m_s = 25$ MeV, and interaction parameter $\alpha_c = 0$.[15] SQM2 has $B = 64.21$ MeV fm^{-3}, $m_s = 150$ MeV, and $\alpha_c = 0.3$, while SQM3 has $B = 57.39$ MeV fm^{-3}, $m_s = 50$ MeV, and $\alpha_c = 0.6$.

Mark Alford provided the equation of state **APR**, which is also taken from Akmal, Pandharipande, and Ravenhall [3, 4]. At high densities, it matches AP4 perfectly, but at low densities, the two do not match. This is because APR features a phase transition, possibly due to neutral pion condensation, between low-density and high-density phases. It is unclear why the Lattimer AP1–AP4 equations of state do not contain this behavior. Lattimer [160] suggests that perhaps the density discontinuity at the point of transition was averaged away; however, in that case, APR and AP4 should also match at low densities.

All of Alford's other equations of state represent hybrid stars, a mix between normal matter (with an APR equation of state) and quark matter. In theory, the quark matter is in a color-superconducting phase, in which quarks form Cooper pairs in the so-called "color-flavor-locked" (CFL) configuration [7]. In the provided equations of state, $m_s = 180$ MeV, and the CFL energy gap $\Delta = 0$. The latter condition means that technically the equations of state are calculated for *unpaired* quark matter. Alford points out [5] that a key result of [6] is that the equation of state depends mainly on the quantity $m_s^2 - 4\Delta^2$, so that this discrepancy may not matter. However, in the future, it would be nice to use a "true" color-flavor-locked EOS with $\Delta > 0$ and compare the results.

Alford's hybrid equations of state are:

- **ALF2N**: The bag constant is adjusted so that the transition between nuclear matter and quark matter occurs at *twice* the nuclear density. The quarks are considered to be free and noninteracting; in the language of [6], the parameter

[15] These parameters, given in the data file, do not match the description of SQM1 given in [161]. Those parameters are: $B = 94.92$ MeV fm^{-3}, $m_s = 0$, and $\alpha_c = 0$. Jim Lattimer [160] notes that the difference in strange mass should not have a significant effect; in any case, the mass-radius relation matches that of [161].

$c = 0$. Finally, the two regions (nuclear and quark matter) do not mix ("**N**o mixing"); instead, there is a sharp transition with a density discontinuity.

- **ALF2M**: Same as ALF2N, but the transition between the nuclear and quark matter phases is smoothed out to avoid a density discontinuity. (The phases are "**M**ixed".)

- **ALF3N**: Same as ALF2N, but the transition occurs at *three* times nuclear density.

- **ALF3M**: Same as ALF3N, but with a smooth transition.

- **ALF3C**: The quark matter parameters of this equation of state are the same as ALF3N and ALF3M, except that perturbative QCD corrections are included. In the language of [6], $c = 0.3$. In addition, this equation of state has quark matter at low densities (much like the SQM stars), transitions to APR nuclear matter at medium densities, and then transitions back to quark matter at high densities. Multiple transitions like this are possible because $c = 0.3$ quark matter has a very similar equation of state to APR nuclear matter.

- **ALF45**: Like ALF3C, this equation of state also features QCD interactions ($c = 0.3$), but like the others, it only has a single transition. It occurs at 4.5 times nuclear density and is smooth.

Figure 4-4 shows the complete equation of state in all cases which use the FMT, BPS, BBP, and NV functions at low density (that is, all but SQM1-SQM3 and ALF3C). Figure 4-5 zooms in on high densities for Lattimer equations of state which contain only nucleons. We see that all of these equations of state have somewhat similar slopes (hence the usual polytrope approximation), although the absolute value of pressure can vary by ~ 5. Despite the similarities, each equation generates a strikingly unique mass-radius relation, also pictured in the figure. Lattimer and Prakash [161] point out that stiffer equations of state, such as MS1 and MS3, produce the largest masses; meanwhile, softer equations of state, such as MS2, PAL6, AP1, AP2,

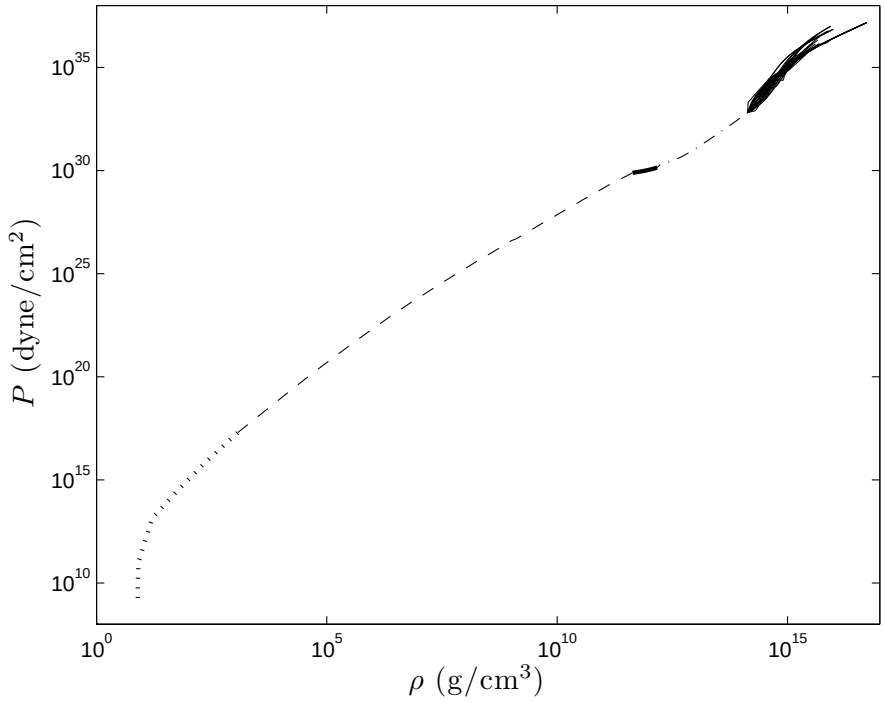

Figure 4-4: Complete equation of state for cases with normal matter at the surface. Starting from low densities, the various pieces are: FMT (dotted line), BPS (dashed line), BBP (thick solid line), NV (dot-dashed line), and various high density EOS (thin solid lines). The following figures provide close-ups on these high density equations of state, as well as those including quark matter which do not have this low-density behavior.

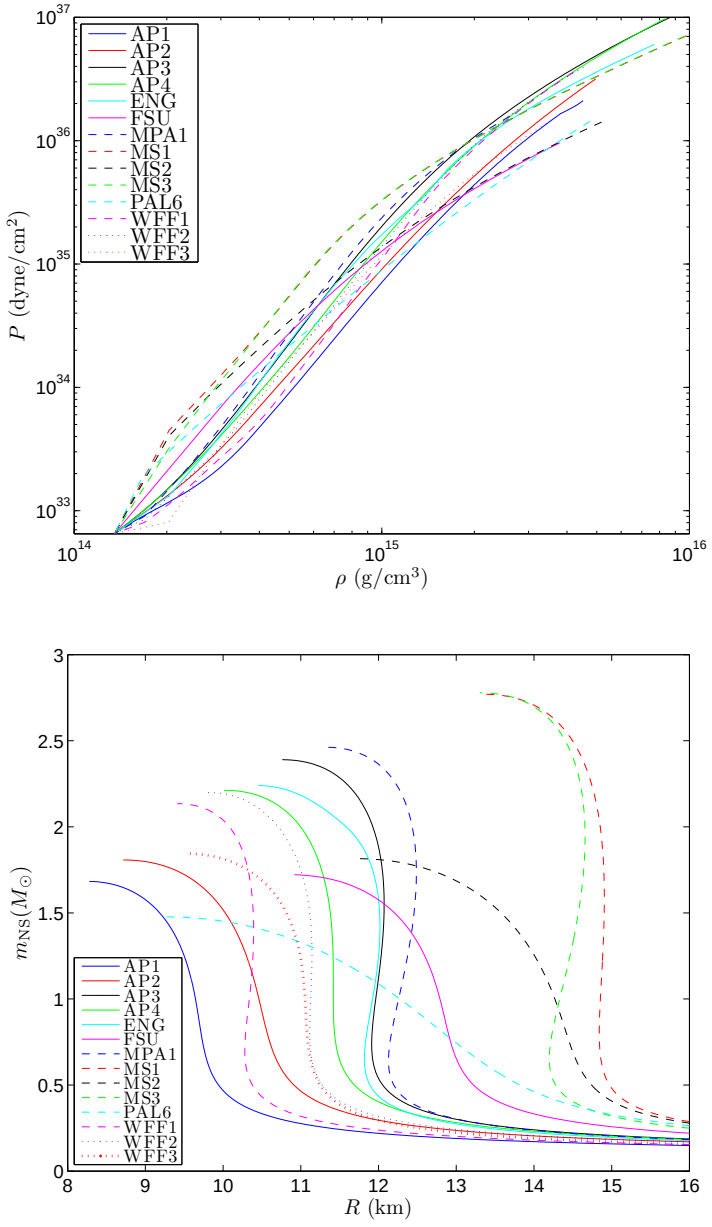

Figure 4-5: Top: High density equations of state provided by Jim Lattimer and which include only nucleons. Bottom: Mass-radius relations for these equations of state.

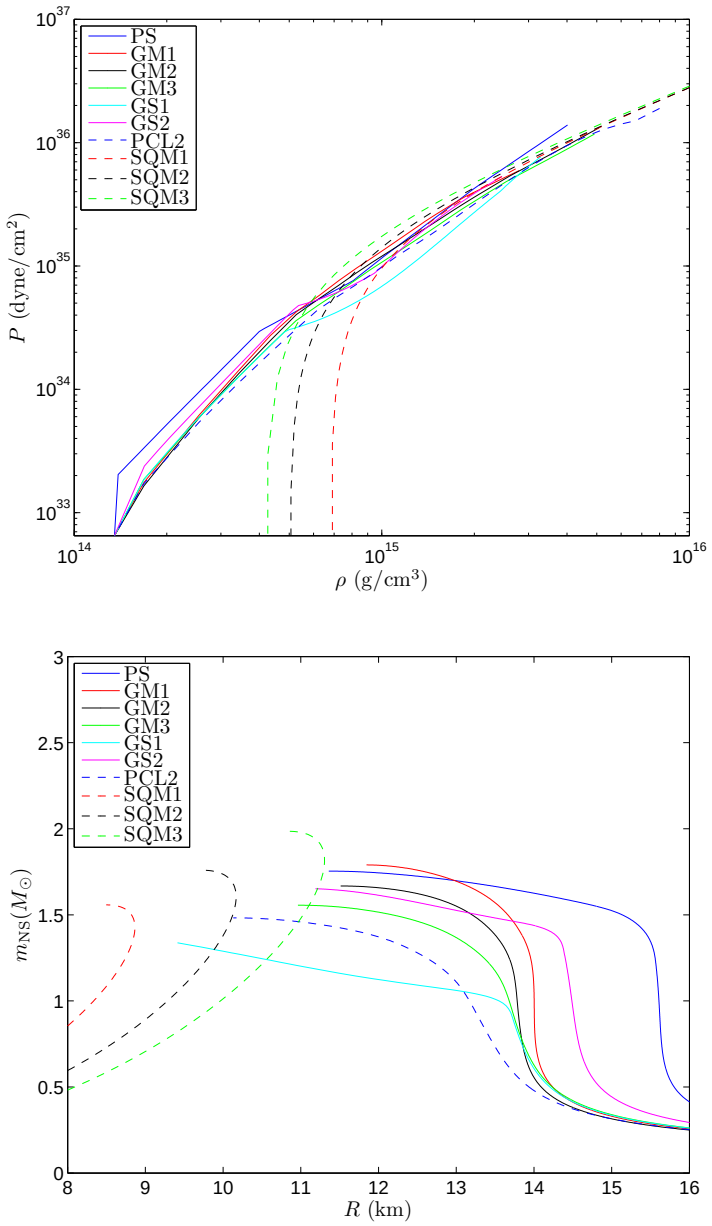

Figure 4-6: Top: High density equations of state provided by Jim Lattimer and which include more exotic components, including pions (PS), hyperons (GM1-GM3), kaons (GS1-GS2), hyperons and free quarks (PCL2), and strange quark matter (SQM1-SQM3). While most of these equations of state match to the low density EOS pictured in Fig. 4-4, the SQM EOS reach zero pressure at a large finite density. Bottom: Mass-radius relations for these equations of state.

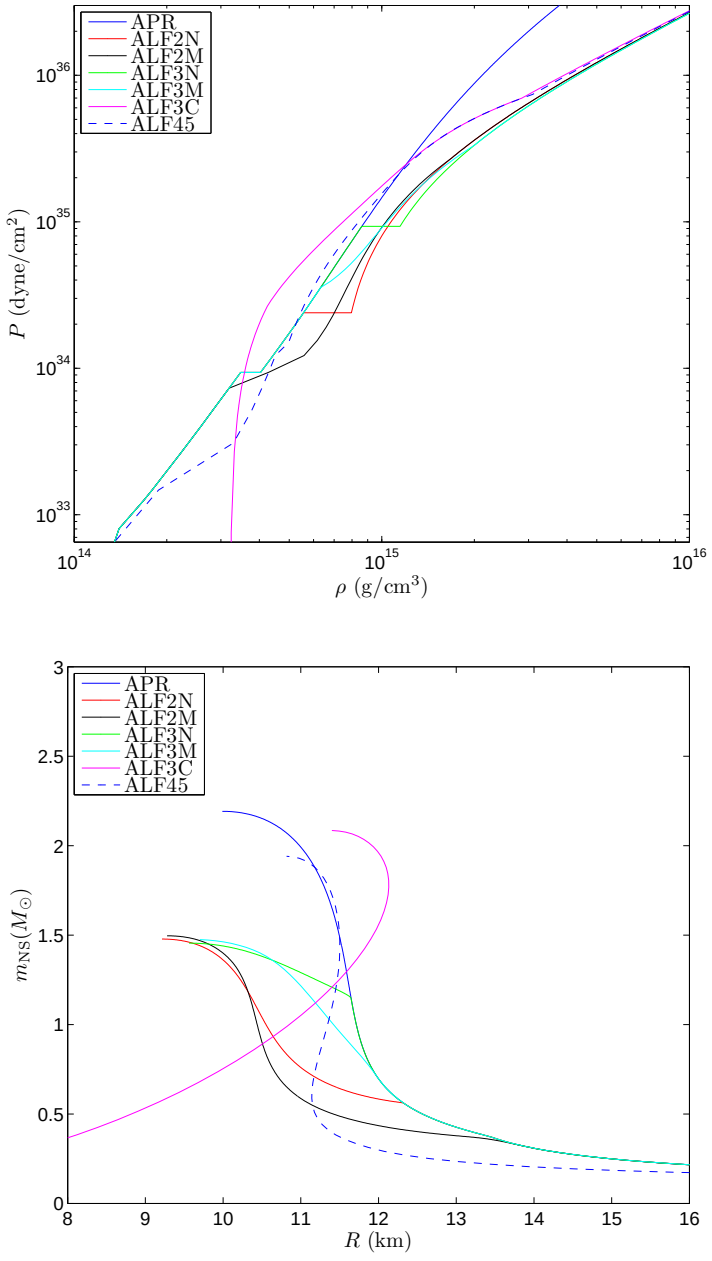

Figure 4-7: Top: High density equations of state provided by Mark Alford. APR includes only nucleons (except for a pion condensation phase transition), while all the others include quark matter as described in the text. All but ALF3C match to the low density EOS pictured in Fig. 4-4. Bottom: Mass-radius relations for these equations of state.

and FSU, produce smaller masses. The radius is controlled mainly by the pressure around nuclear density; higher pressures equal larger radii. The dependence of radius on mass is stronger for softer equations of state.

Figure 4-6 pictures the high density equations of state and mass-radius relations when more exotic components (pions, kaons, hyperons, quarks) are included in the analysis. The exotic components generally soften the equation of state, resulting in smaller masses and more dependence of radius on mass. The SQM stars are in their own class. Because they are self-bound, they have no minimum mass; in addition, radius generally increases, rather than decreases, with mass.

Finally, Figure 4-7 pictures the equations of state taken from Mark Alford, as well as their mass-radius relationships. All of the various transitions can be seen clearly. Quark matter in general softens the equations of state; however, those with QCD interactions (ALF3C and ALF45) look more like normal matter and thus have higher maximum masses. ALF3C, because it transitions back to quark matter at low densities, has a mass-radius profile similar to the SQM stars.

4.7.3 General equations of state: Love numbers

We now turn to the key result of this section, the Love numbers for different realistic equations of state. Figure 4-8 presents k_2 for stars containing only nucleons. Note that the shape is quite different from that of Figure 4-3. In the polytrope case, the polytrope equation of state is considered to hold for all densities down to zero. In reality, at lower densities the equation of state will change to the various forms described above. Stars with low $m_{\rm NS}/R$ are described by these equations of state throughout more of their structure. These equations of state are softer than the high density EOS, and the Love number is consequently quite small.

In the range of interest, $0.1 \lesssim m_{\rm NS}/R \lesssim 0.2$, the Love numbers are close to the polytrope value. The polytrope is larger than most (but not all) equations of state for smaller $m_{\rm NS}/R$ but eventually drops below all of them at higher $m_{\rm NS}/R$. In general, the Love numbers for realistic equations of state are remarkably similar, as seen by the crowded nature of the plot (which is both a flaw and a feature!). Very soft equations

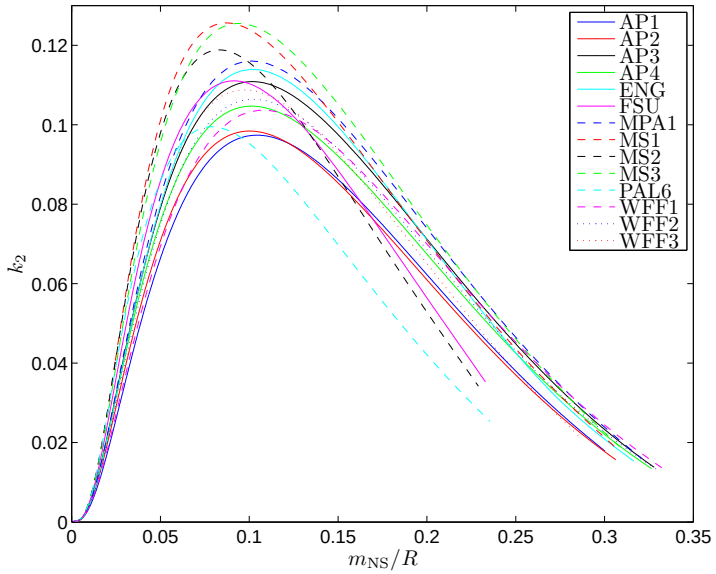

Figure 4-8: Love number k_2 as a function of compactness $m_{\rm NS}/R$ for Lattimer equations of state which include only nucleons.

of state like PAL6, AP1, and AP2 produce the smallest response; others, like MS2 and FSU, are quite small at high $m_{\rm NS}/R$ but among the the highest at $m_{\rm NS}/R \sim 0.1$.

Figure 4-9 shows the Love numbers for equations of state including exotic components, with the exception of the SQM stars which have much different results. The Love numbers at small $m_{\rm NS}/R$ resemble those for the nucleon-only case, but the results at $m_{\rm NS}/R \sim 0.2$ are a bit smaller, as expected since the exotic components tend to soften the equation of state.

Figure 4-10 shows the results for Alford's hybrid equations of state, with the exception of ALF3C. These stars have smaller Love numbers than the nucleon-only average both at high and low $m_{\rm NS}/R$. The exception is ALF45; it seems the QCD corrections stiffen the equation of state and increase the Love number.

Finally, Figure 4-11 depicts the Love numbers for stars composed entirely of quarks, SQM1, SQM2, SQM3, and ALF3C. These results were held off to a separate plot for reasons of scale. Because the all-quark equations of state drop suddenly to zero pressure at large finite density, without including the standard low-density

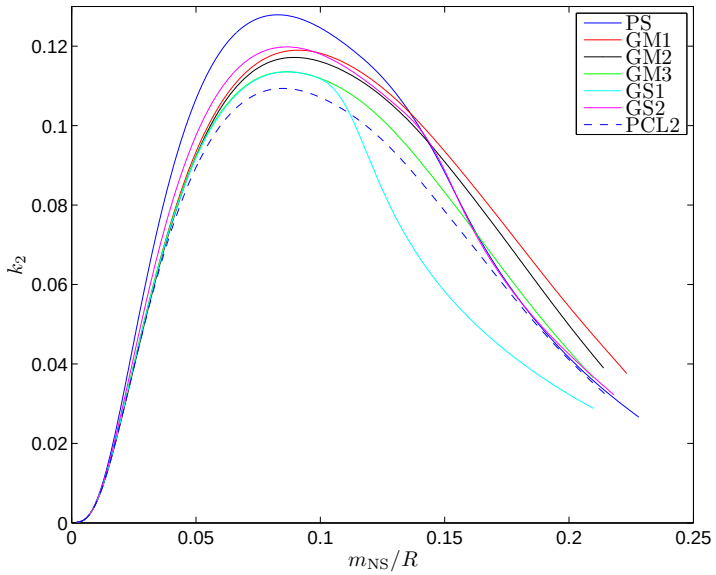

Figure 4-9: Love number k_2 as a function of compactness $m_{\rm NS}/R$ for Lattimer equations of state which include exotic components (excluding SQM1-SQM3).

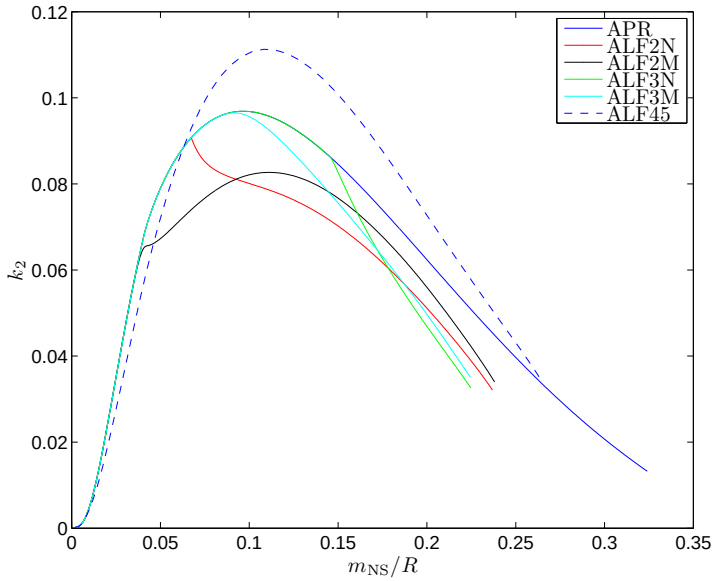

Figure 4-10: Love number k_2 as a function of compactness $m_{\rm NS}/R$ for equations of state provided by Mark Alford (excluding ALF3C).

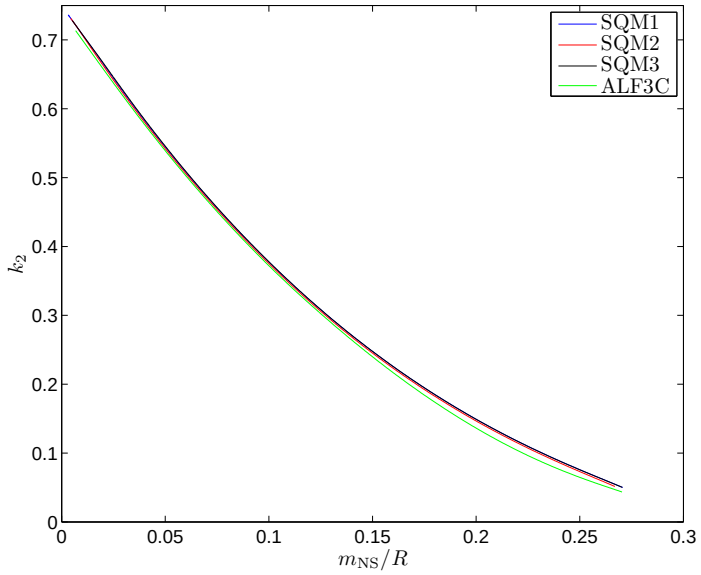

Figure 4-11: Love number k_2 as a function of compactness $m_{\rm NS}/R$ for equations of state SQM1-SQM3 and ALF3C.

behavior, they effectively behave like an $n = 0$ ($\gamma = \infty$) polytrope as $m_{\rm NS}/R \to 0$. Indeed, the curves asymptote to $k_2 = 0.75$, the $n = 0$ Newtonian Love number given in [40]. As a result of this behavior, the all-quark models feature Love numbers which can be quite large in the region of interest. They are also the only ones substantially different from the polytrope results.

To quantify the impact of these Love numbers on the observed GWs, we calculate the number of additional gravitational wave cycles $\Delta N_{\rm cyc}$. To do so, we use (4.120) with $f_{\rm low} = 10$ Hz and $f_{\rm high} = 1000$ Hz. These values are appropriate for Advanced LIGO. Note that Flanagan and Hinderer [95] only trust their analysis up to 400 Hz. However, the tidal effect scales as $f^{5/3}$ and is thus more important at higher frequencies. By including frequencies up to 1000 Hz, we are technically extending the result past its region of validity in order to probe (albeit with some error) a more relevant regime. It would be worthwhile to investigate this tradeoff further.

Table 4.3 shows the results for neutron star binaries in which each star has $m_{\rm NS} = 1.4\ M_\odot$. (See the caption for one exception.) We see that for the most part, the tidal

Equation of state	$\Delta N_{\rm cyc}$
AP1	-0.0673
AP2	-0.117
AP3	-0.406
AP4	-0.272
ENG	-0.402
FSU	-0.418
MPA1	-0.505
MS1	-1.44
MS2	-0.791
MS3	-1.30
PAL6	-0.104
WFF1	-0.157
WFF2	-0.239
WFF3	-0.201
PS	-1.68
GM1	-0.867
GM2	-0.715
GM3	-0.480
GS1	-0.0413
GS2	-1.00
PCL2	-0.241
SQM1	-0.0981
SQM2	-0.253
SQM3	-0.462
APR	-0.270
ALF2N	-0.0808
ALF2M	-0.0958
ALF3N	-0.104
ALF3M	-0.118
ALF3C	-0.761
ALF45	-0.313

Table 4.3: Additional number of GW cycles $\Delta N_{\rm cyc}$ caused by tidal effects for NS-NS binaries with $m_1 = m_2 = 1.4 M_\odot$. (The exception is GS1, which has a maximum mass less than $1.4 M_\odot$. In that case, each neutron star has the maximum mass.)

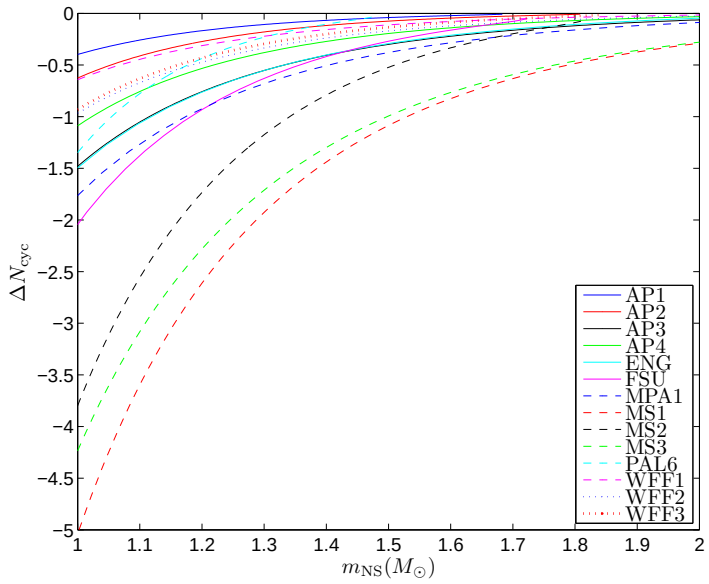

Figure 4-12: Number of GW cycles $\Delta N_{\rm cyc}$ contributed by tidal term for Lattimer equations of state which include only nucleons. Each neutron star in the binary has mass $m_{\rm NS}$.

terms affect the phase very little, adjusting it by less than 1 cycle. The only equations of state which do better than 1 cycle are MS1, MS3, PS, and GS2. Looking back at the mass-radius relations, we see that these equations of state have $R \sim 15 - 16$ km at $m_{\rm NS} = 1.4\ M_\odot$. Since $\Delta N_{\rm cyc} \sim \lambda \sim R^5$, the large radii of these stars is the primary cause of this effect. Meanwhile, the quark stars, SQM1-SQM3 and ALF3C, have large Love numbers but moderate to small radii, so the effect is not as large.

We see in (4.120) that the tidal effect scales inversely with mass. Figures 4-12, 4-13, and 4-14 show $\Delta N_{\rm cyc}$ for stars in the mass range $1-2\ M_\odot$. At high mass, the effect is often very small. (In the first plot, this makes it hard to distinguish these values from each other, but again, this fault of the plot is also a *feature*: $\Delta N_{\rm cyc}$ is so small that it is essentially irrelevant. Certainly it cannot be used to distinguish equations of state from one another.) At low mass, the effect grows. For a neutron star of mass 1.2 $M_\odot$, the smallest which has been reliably measured, several equations of state produce effects of greater than 1 cycle, including a majority of Lattimer's exotic EOS. Some equations of state produce an effect of several cycles.

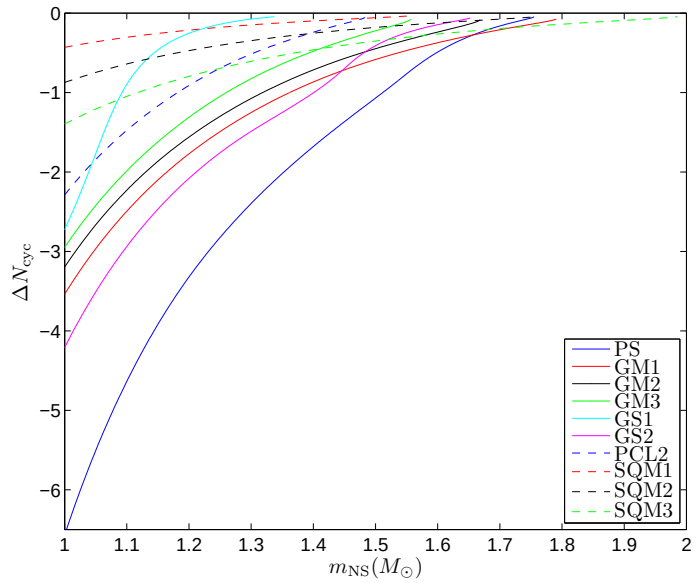

Figure 4-13: Number of GW cycles $\Delta N_{\rm cyc}$ contributed by tidal term for Lattimer equations of state which include exotic components. Each neutron star in the binary has mass $m_{\rm NS}$.

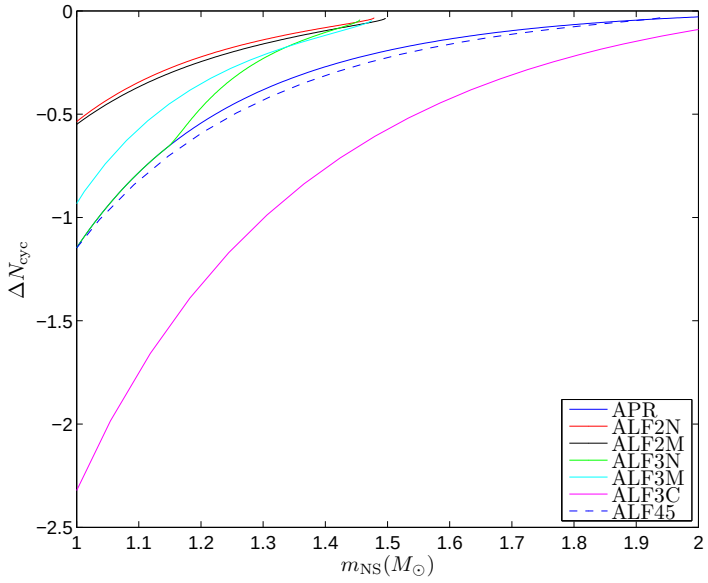

Figure 4-14: Number of GW cycles $\Delta N_{\rm cyc}$ contributed by tidal term for equations of state provided by Mark Alford. Each neutron star in the binary has mass $m_{\rm NS}$.

4.8 Summary and conclusions

The equation of state at the center of neutron stars is poorly understood. The measurement of tidal effects in binary systems may help distinguish between the theoretical possibilities. Because neutron stars are so compact ($m_{\rm NS}/R \sim 0.1 - 0.2$), these effects are properly calculated using relativistic self-gravity. Our approach is to use the $\ell = 2$ nonradial perturbations to the standard TOV stellar structure solution. In this chapter, we have used this formalism to study two distinct tidal effects.

First, we have investigated the equilibrium structure of a neutron star sitting in the *specific* tidal field of a black hole. At some orbital radius $r_{\rm tide}$, no equilibrium solution exists, and the star is tidally disrupted by the hole. This disruption will be seen in the gravitational wave signal from the binary, at the frequency corresponding to $r_{\rm tide}$. By using relativistic self-gravity for the neutron star, we find that $r_{\rm tide}$ is much smaller than in the Newtonian self-gravity limit. The star's stronger gravity helps it to avoid early disruption. Unfortunately, when we compare our results to those generated by numerical relativity, they are not as accurate as hoped. In addition to being a perturbative (and thus intrinsically weak-field) procedure, our method neglects the bulk rotation of the star and any internal fluid motions, both of which can have a significant effect on the distorted structure. It would be interesting in the future to incorporate both tidal and rotational distortions into this work and see how close we can come to the numerical relativity results. However, doing so would require reformulating the perturbation theory and is beyond the scope of this thesis.

Even without including these effects, our method may still be an excellent way to study the tidal deformation of neutron stars at large orbital separation. One possible application of this work is the same as Alvi's initial inspiration: the generation of initial data for numerical relativity simulations. To do this, we would need to perform a full match between the perturbed neutron star spacetime and the near zone post-Newtonian metric. This would require a deeper understanding of Alvi's work, including the mistakes he made [261]. It may even be prudent to do a more formal, less *ad hoc* matching, as suggested by Poisson [193]. Such a project is also beyond

the scope of this thesis.

We have also studied the excitation of neutron star modes in a neutron star-neutron star binary. These excitations, and their effect on the gravitational wave phase, can be characterized to lowest order by a single quantity known as the Love number. The Love number can be described as the neutron star's response to a *general* (static) tidal field. By using pieces of our disruption code, we have determined Love numbers for a wide variety of realistic equations of state and then translated these results into the change in accumulated gravitational wave cycles in the Advanced LIGO band. For neutron stars of $\sim 1.4\ M_\odot$, most equations of state contribute less than one cycle to the GW signal; however, a few contribute more than this. The effect decreases with mass, so a low-mass binary with the right equation of state could be reasonably impacted by the tidal terms. Further work is necessary to carefully characterize the ability of Advanced LIGO to measure the Love number. In particular, a Fisher matrix analysis should be performed, as in Chapters 2 and 3.

Another interesting question is whether the tidal effects could impact *detection* of the gravitational waves. Even though the tidal term is formally of 5PN order, it can become as important or more important than 3PN and 3.5PN point-particle terms for low-mass stars [34]. When $\Delta N_{\rm cyc} \gtrsim 1$ for any post-Newtonian term, matched filtering templates which neglect that term may fail to detect the signal. It would be good to investigate this question further.

In the end, of course, all questions about tidal disruption and distortion can be answered by numerical relativity. However, while those simulations remain difficult and computationally expensive, it is important to consider approximate approaches like these in order to prepare for the coming age of gravitational wave science.

Bibliography

[1] A. Abramovici, W. E. Althouse, R. W. P. Drever, Y. Gürsel, S. Kawamura, F. J. Raab, D. Shoemaker, L. Sievers, R. E. Spero, K. S. Thorne, R. E. Vogt, R. Weiss, S. E. Whitcomb, and M. E. Zucker. LIGO: The Laser Interferometer Gravitational-Wave Observatory. *Science*, 256:325–333, April 1992.

[2] M. L. Aizenman. The equilibrium and the stability of the Roche-Riemann ellipsoids. *The Astrophysical Journal*, 153:511–544, August 1968.

[3] A. Akmal and V. R. Pandharipande. Spin-isospin structure and pion condensation in nucleon matter. *Physical Review C*, 56:2261–2279, October 1997.

[4] A. Akmal, V. R. Pandharipande, and D. G. Ravenhall. Equation of state of nucleon matter and neutron star structure. *Physical Review C*, 58:1804–1828, September 1998.

[5] M. Alford. Private communication, April 2009.

[6] M. Alford, M. Braby, M. Paris, and S. Reddy. Hybrid stars that masquerade as neutron stars. *The Astrophysical Journal*, 629:969–978, August 2005.

[7] M. G. Alford, A. Schmitt, K. Rajagopal, and T. Schäfer. Color superconductivity in dense quark matter. *Reviews of Modern Physics*, 80:1455–1515, October 2008.

[8] K. Alvi. Approximate binary-black-hole metric. *Physical Review D*, 61:124013, June 2000.

[9] M. Ando and the TAMA collaboration. Current status of TAMA. *Classical and Quantum Gravity*, 19:1409–1419, April 2002.

[10] T. A. Apostolatos, C. Cutler, G. J. Sussman, and K. S. Thorne. Spin-induced orbital precession and its modulation of the gravitational waveforms from merging binaries. *Physical Review D*, 49:6274–6297, June 1994.

[11] P. J. Armitage and P. Natarajan. Accretion during the merger of supermassive black holes. *The Astrophysical Journal*, 567:L9–L12, March 2002.

[12] P. J. Armitage and P. Natarajan. Eccentricity of supermassive black hole binaries coalescing from gas-rich mergers. *The Astrophysical Journal*, 634:921–927, December 2005.

[13] J. W. Armstrong, F. B. Estabrook, and M. Tinto. Time-delay interferometry for space-based gravitational wave searches. *The Astrophysical Journal*, 527:814–826, December 1999.

[14] W. D. Arnett and R. L. Bowers. A microscopic interpretation of neutron star structure. *The Astrophysical Journal Supplement Series*, 33:415–436, April 1977.

[15] K. G. Arun, S. Babak, E. Berti, N. Cornish, C. Cutler, J. Gair, S. A. Hughes, B. R. Iyer, R. N. Lang, I. Mandel, E. K. Porter, B. S. Sathyaprakash, S. Sinha, A. M. Sintes, M. Trias, C. Van Den Broeck, and M. Volonteri. Massive black-hole binary inspirals: results from the LISA parameter estimation taskforce. *Classical and Quantum Gravity*, 26:094027, May 2009.

[16] K. G. Arun, B. R. Iyer, B. S. Sathyaprakash, S. Sinha, and C. van den Broeck. Higher signal harmonics, LISA's angular resolution, and dark energy. *Physical Review D*, 76:104016, November 2007.

[17] W. Baade and F. Zwicky. Cosmic rays from super-novae. *Proceedings of the National Academy of Sciences of the United States of America*, 20:259–263, May 1934.

[18] S. Babak, M. Hannam, S. Husa, and B. Schutz. Resolving super massive black holes with LISA. arXiv:0806.1591, June 2008.

[19] J. G. Baker, J. Centrella, D.-I. Choi, M. Koppitz, and J. van Meter. Gravitational-wave extraction from an inspiraling configuration of merging black holes. *Physical Review Letters*, 96:111102, March 2006.

[20] J. G. Baker, S. T. McWilliams, J. R. van Meter, J. Centrella, D.-I. Choi, B. J. Kelly, and M. Koppitz. Binary black hole late inspiral: Simulations for gravitational wave observations. *Physical Review D*, 75:124024, June 2007.

[21] L. Barack and C. Cutler. LISA capture sources: Approximate waveforms, signal-to-noise ratios, and parameter estimation accuracy. *Physical Review D*, 69:082005, April 2004.

[22] G. Baym, H. A. Bethe, and C. J. Pethick. Neutron star matter. *Nuclear Physics A*, 175:225–271, November 1971.

[23] G. Baym, C. Pethick, and P. Sutherland. The ground state of matter at high densities: Equation of state and stellar models. *The Astrophysical Journal*, 170:299–317, December 1971.

[24] M. C. Begelman, R. D. Blandford, and M. J. Rees. Massive black hole binaries in active galactic nuclei. *Nature*, 287:307–309, September 1980.

[25] K. Belczynski, R. E. Taam, V. Kalogera, F. A. Rasio, and T. Bulik. On the rarity of double black hole binaries: Consequences for gravitational wave detection. *The Astrophysical Journal*, 662:504–511, June 2007.

[26] P. Bender, A. Brillet, I. Ciufolini, A. M. Cruise, C. Cutler, K. Danzmann, F. Fidecaro, W. M. Folkner, J. Hough, P. McNamara, M. Peterseim, D. Robertson, M. Rodrigues, A. Rüdiger, M. Sandford, G. Schäfer, R. Schilling, B. Schutz, C. Speake, R. T. Stebbins, T. Sumner, P. Touboul, J.-Y. Vinet, S. Vitale, H. Ward, and W. Winkler. LISA Pre-Phase A Report, Second Edition, July 1998.

[27] P. Berczik, D. Merritt, R. Spurzem, and H.-P. Bischof. Efficient merger of binary supermassive black holes in nonaxisymmetric galaxies. *The Astrophysical Journal*, 642:L21–L24, May 2006.

[28] I. Berentzen, M. Preto, P. Berczik, D. Merritt, and R. Spurzem. Binary black hole merger in galactic nuclei: Post-Newtonian simulations. *The Astrophysical Journal*, 695:455–468, April 2009.

[29] E. Berti, A. Buonanno, and C. M. Will. Estimating spinning binary parameters and testing alternative theories of gravity with LISA. *Physical Review D*, 71:084025, April 2005.

[30] E. Berti, V. Cardoso, and C. M. Will. Gravitational-wave spectroscopy of massive black holes with the space interferometer LISA. *Physical Review D*, 73:064030, March 2006.

[31] E. Berti, S. Iyer, and C. M. Will. Post-Newtonian diagnosis of quasiequilibrium configurations of neutron star-neutron star and neutron star-black hole binaries. *Physical Review D*, 77:024019, January 2008.

[32] S. Bianchi, M. Chiaberge, E. Piconcelli, M. Guainazzi, and G. Matt. Chandra unveils a binary active galactic nucleus in Mrk 463. *Monthly Notices of the Royal Astronomical Society*, 386:105–110, May 2008.

[33] L. Bildsten and C. Cutler. Tidal interactions of inspiraling compact binaries. *The Astrophysical Journal*, 400:175–180, November 1992.

[34] L. Blanchet. Gravitational radiation from post-Newtonian sources and inspiralling compact binaries. *Living Reviews in Relativity*, 9:4, June 2006.

[35] L. Blanchet, A. Buonanno, and G. Faye. Higher-order spin effects in the dynamics of compact binaries. II. Radiation field. *Physical Review D*, 74:104034, November 2006.

[36] L. Blanchet, T. Damour, B. R. Iyer, C. M. Will, and A. G. Wiseman. Gravitational-radiation damping of compact binary systems to second post-Newtonian order. *Physical Review Letters*, 74:3515–3518, May 1995.

[37] R. D. Blandford and R. L. Znajek. Electromagnetic extraction of energy from Kerr black holes. *Monthly Notices of the Royal Astronomical Society*, 179:433–456, May 1977.

[38] J. N. Bode and E. Phinney. Observability of circumbinary disks following massive black hole mergers. In *Bulletin of the American Astronomical Society*, volume 41, page 341, January 2009.

[39] C. Bradaschia, R. del Fabbro, A. di Virgilio, A. Giazotto, H. Kautzky, V. Montelatici, D. Passuello, A. Brillet, O. Cregut, P. Hello, C. N. Man, P. T. Manh, A. Marraud, D. Shoemaker, J. Y. Vinet, F. Barone, L. di Fiore, L. Milano, G. Russo, J. M. Aguirregabiria, H. Bel, J. P. Duruisseau, G. Le Denmat, P. Tourrenc, M. Capozzi, M. Longo, M. Lops, I. Pinto, G. Rotoli, T. Damour, S. Bonazzola, J. A. Marck, Y. Gourghoulon, L. E. Holloway, F. Fuligni, V. Iafolla, and G. Natale. The VIRGO project: A wide band antenna for gravitational wave detection. *Nuclear Instruments and Methods in Physics Research Section A*, 289:518–525, April 1990.

[40] R. A. Brooker and T. W. Olle. Apsidal-motion constants for polytropic models. *Monthly Notices of the Royal Astronomical Society*, 115:101–106, 1955.

[41] A. Buonanno, Y. Chen, Y. Pan, H. Tagoshi, and M. Vallisneri. Detecting gravitational waves from precessing binaries of spinning compact objects. II. Search implementation for low-mass binaries. *Physical Review D*, 72:084027, October 2005.

[42] A. Buonanno, Y. Chen, Y. Pan, and M. Vallisneri. Quasiphysical family of gravity-wave templates for precessing binaries of spinning compact objects: Application to double-spin precessing binaries. *Physical Review D*, 70:104003, November 2004.

[43] A. Buonanno, Y. Chen, and M. Vallisneri. Detecting gravitational waves from precessing binaries of spinning compact objects: Adiabatic limit. *Physical Review D*, 67:104025, May 2003.

[44] A. Buonanno, G. B. Cook, and F. Pretorius. Inspiral, merger, and ring-down of equal-mass black-hole binaries. *Physical Review D*, 75:124018, June 2007.

[45] M. Campanelli, C. O. Lousto, P. Marronetti, and Y. Zlochower. Accurate evolutions of orbiting black-hole binaries without excision. *Physical Review Letters*, 96:111101, March 2006.

[46] M. Campanelli, C. O. Lousto, Y. Zlochower, and D. Merritt. Maximum gravitational recoil. *Physical Review Letters*, 98:231102, June 2007.

[47] B. W. Carroll and D. A. Ostlie. *An Introduction to Modern Astrophysics*. Addison Wesley, 1996.

[48] S. M. Carroll. *Spacetime and Geometry: An Introduction to General Relativity*. Addison Wesley, 2004.

[49] S. Chandrasekhar. The maximum mass of ideal white dwarfs. *The Astrophysical Journal*, 74:81–82, July 1931.

[50] S. Chandrasekhar. *Ellipsoidal Figures of Equilibrium*. Dover Publications, Inc., 1987.

[51] N. Christensen, R. J. Dupuis, G. Woan, and R. Meyer. Metropolis-Hastings algorithm for extracting periodic gravitational wave signals from laser interferometric detector data. *Physical Review D*, 70:022001, July 2004.

[52] N. Christensen and R. Meyer. Using Markov chain Monte Carlo methods for estimating parameters with gravitational radiation data. *Physical Review D*, 64:022001, July 2001.

[53] E. J. M. Colbert and M. C. Miller. Observational evidence for intermediate-mass black holes in ultra-luminous X-ray sources. In M. Novello, S. Perez Bergliaffa, and R. Ruffini, editors, *The Tenth Marcel Grossmann Meeting. On recent developments in theoretical and experimental general relativity, gravitation and relativistic field theories*, page 530, January 2005.

[54] The LIGO Scientific Collaboration and K. C. Hurley. Implications for the origin of GRB 070201 from LIGO observations. *The Astrophysical Journal*, 681:1419–1430, July 2008.

[55] N. J. Cornish. Gravitational wave confusion noise. arXiv:gr-qc/0304020, April 2003.

[56] N. J. Cornish. Protected observing periods for LISA. LISA International Science Team white paper, available at `http://gravity.psu.edu/~shane/pdf/njcProtected.pdf`, October 2005.

[57] N. J. Cornish. Private communication, September 2007.

[58] N. J. Cornish and J. Crowder. LISA data analysis using Markov chain Monte Carlo methods. *Physical Review D*, 72:043005, August 2005.

[59] N. J. Cornish and E. K. Porter. MCMC exploration of supermassive black hole binary inspirals. *Classical and Quantum Gravity*, 23:S761–S767, October 2006.

[60] N. J. Cornish and E. K. Porter. Catching supermassive black hole binaries without a net. *Physical Review D*, 75:021301, January 2007.

[61] N. J. Cornish and L. J. Rubbo. LISA response function. *Physical Review D*, 67:022001, January 2003.

[62] C. Cutler. Angular resolution of the LISA gravitational wave detector. *Physical Review D*, 57:7089–7102, June 1998.

[63] C. Cutler and É. É. Flanagan. Gravitational waves from merging compact binaries: How accurately can one extract the binary's parameters from the inspiral waveform? *Physical Review D*, 49:2658–2697, March 1994.

[64] N. Dalal, D. E. Holz, X. Chen, and J. A. Frieman. Corrective lenses for high-redshift supernovae. *The Astrophysical Journal*, 585:L11–L14, March 2003.

[65] G. Darmois. Les équations de la gravitation einsteinienne. *Mémorial des Sciences Mathématiques*, 25:30, 1927.

[66] G. H. Darwin. On the figure and stability of a liquid satellite. *Philosophical Transactions of the Royal Society A*, 206:161–248, 1906.

[67] S. Detweiler. Pulsar timing measurements and the search for gravitational waves. *The Astrophysical Journal*, 234:1100–1104, December 1979.

[68] S. Detweiler. Private communication, June 2008.

[69] T. Di Matteo, J. Colberg, V. Springel, L. Hernquist, and D. Sijacki. Direct cosmological simulations of the growth of black holes and galaxies. *The Astrophysical Journal*, 676:33–53, March 2008.

[70] M. Dotti, R. Salvaterra, A. Sesana, M. Colpi, and F. Haardt. On the search of electromagnetic cosmological counterparts to coalescences of massive black hole binaries. *Monthly Notices of the Royal Astronomical Society*, 372:869–875, October 2006.

[71] O. Dreyer, B. Kelly, B. Krishnan, L. S. Finn, D. Garrison, and R. Lopez-Aleman. Black-hole spectroscopy: testing general relativity through gravitational-wave observations. *Classical and Quantum Gravity*, 21:787–803, February 2004.

[72] S. Droz, D. J. Knapp, E. Poisson, and B. J. Owen. Gravitational waves from inspiraling compact binaries: Validity of the stationary-phase approximation to the Fourier transform. *Physical Review D*, 59:124016, June 1999.

[73] F. Echeverria. Gravitational-wave measurements of the mass and angular momentum of a black hole. *Physical Review D*, 40:3194–3203, November 1989.

[74] A. Eckart and R. Genzel. Stellar proper motions in the central 0.1 pc of the Galaxy. *Monthly Notices of the Royal Astronomical Society*, 284:576–598, January 1997.

[75] A. Einstein. Zur Elektrodynamik bewegter Körper. *Annalen der Physik*, 322:891–921, June 1905.

[76] A. Einstein. Die Grundlage der allgemeinen Relativitätstheorie. *Annalen der Physik*, 354:769–822, May 1916.

[77] L. Engvik, E. Osnes, M. Hjorth-Jensen, G. Bao, and E. Ostgaard. Asymmetric nuclear matter and neutron star properties. *The Astrophysical Journal*, 469:794–807, October 1996.

[78] A. Escala, R. B. Larson, P. S. Coppi, and D. Mardones. The role of gas in the merging of massive black holes in galactic nuclei. I. Black hole merging in a spherical gas cloud. *The Astrophysical Journal*, 607:765–777, June 2004.

[79] A. Escala, R. B. Larson, P. S. Coppi, and D. Mardones. The role of gas in the merging of massive black holes in galactic nuclei. II. Black hole merging in a nuclear gas disk. *The Astrophysical Journal*, 630:152–166, September 2005.

[80] Z. B. Etienne, J. A. Faber, Y. T. Liu, S. L. Shapiro, K. Taniguchi, and T. W. Baumgarte. Fully general relativistic simulations of black hole-neutron star mergers. *Physical Review D*, 77:084002, April 2008.

[81] Z. B. Etienne, Y. T. Liu, S. L. Shapiro, and T. W. Baumgarte. General relativistic simulations of black-hole-neutron-star mergers: Effects of black-hole spin. *Physical Review D*, 79:044024, February 2009.

[82] J. A. Faber, P. Grandclément, F. A. Rasio, and K. Taniguchi. Measuring neutron-star radii with gravitational-wave detectors. *Physical Review Letters*, 89:231102, November 2002.

[83] X. Fan, V. K. Narayanan, R. H. Lupton, M. A. Strauss, G. R. Knapp, R. H. Becker, R. L. White, L. Pentericci, S. K. Leggett, Z. Haiman, J. E. Gunn, Ž. Ivezić, D. P. Schneider, S. F. Anderson, J. Brinkmann, N. A. Bahcall, A. J. Connolly, I. Csabai, M. Doi, M. Fukugita, T. Geballe, E. K. Grebel, D. Harbeck, G. Hennessy, D. Q. Lamb, G. Miknaitis, J. A. Munn, R. Nichol, S. Okamura, J. R. Pier, F. Prada, G. T. Richards, A. Szalay, and D. G. York. A survey of $z > 5.8$ quasars in the Sloan Digital Sky Survey. I. Discovery of three new quasars and the spatial density of luminous quasars at $z \sim 6$. *The Astronomical Journal*, 122:2833–2849, December 2001.

[84] A. J. Farmer and E. S. Phinney. The gravitational wave background from cosmological compact binaries. *Monthly Notices of the Royal Astronomical Society*, 346:1197–1214, December 2003.

[85] A. J. Faulkner, M. Kramer, A. G. Lyne, R. N. Manchester, M. A. McLaughlin, I. H. Stairs, G. Hobbs, A. Possenti, D. R. Lorimer, N. D'Amico, F. Camilo, and M. Burgay. PSR J1756-2251: A new relativistic double neutron star system. *The Astrophysical Journal*, 618:L119–L122, January 2005.

[86] G. Faye, L. Blanchet, and A. Buonanno. Higher-order spin effects in the dynamics of compact binaries. I. Equations of motion. *Physical Review D*, 74:104033, November 2006.

[87] L. Ferrarese and D. Merritt. A fundamental relation between supermassive black holes and their host galaxies. *The Astrophysical Journal*, 539:L9–L12, August 2000.

[88] V. Ferrari, L. Gualtieri, and F. Pannarale. A semi-relativistic model for tidal interactions in BH-NS coalescing binaries. arXiv:0801.2911v2, February 2009.

[89] R. P. Feynman, N. Metropolis, and E. Teller. Equations of state of elements based on the generalized Fermi-Thomas theory. *Physical Review*, 75:1561–1573, May 1949.

[90] L. S. Finn. Detection, measurement, and gravitational radiation. *Physical Review D*, 46:5236–5249, December 1992.

[91] L. G. Fishbone. The relativistic Roche problem. I. Equilibrium theory for a body in equatorial, circular orbit around a Kerr black hole. *The Astrophysical Journal*, 185:43–68, October 1973.

[92] M. J. Fitchett. The influence of gravitational wave momentum losses on the centre of mass motion of a Newtonian binary system. *Monthly Notices of the Royal Astronomical Society*, 203:1049–1062, June 1983.

[93] É. É. Flanagan. General-relativistic coupling between orbital motion and internal degrees of freedom for inspiraling binary neutron stars. *Physical Review D*, 58:124030, December 1998.

[94] É. É Flanagan. Private communication, February 2009.

[95] É. É. Flanagan and T. Hinderer. Constraining neutron-star tidal Love numbers with gravitational-wave detectors. *Physical Review D*, 77:021502, January 2008.

[96] É. É. Flanagan and S. A. Hughes. Measuring gravitational waves from binary black hole coalescences. I. Signal to noise for inspiral, merger, and ringdown. *Physical Review D*, 57:4535–4565, April 1998.

[97] É. É. Flanagan and S. A. Hughes. The basics of gravitational wave theory. *New Journal of Physics*, 7:204, September 2005.

[98] J. A. Frieman. Weak lensing and the measurement of q_0 from Type Ia supernovae. *Comments on Astrophysics*, 18:323, 1997.

[99] K. Gebhardt, R. Bender, G. Bower, A. Dressler, S. M. Faber, A. V. Filippenko, R. Green, C. Grillmair, L. C. Ho, J. Kormendy, T. R. Lauer, J. Magorrian, J. Pinkney, D. Richstone, and S. Tremaine. A relationship between nuclear black hole mass and galaxy velocity dispersion. *The Astrophysical Journal*, 539:L13–L16, August 2000.

[100] A. M. Ghez, B. L. Klein, M. Morris, and E. E. Becklin. High proper-motion stars in the vicinity of Sagittarius A*: Evidence for a supermassive black hole at the center of our Galaxy. *The Astrophysical Journal*, 509:678–686, December 1998.

[101] N. K. Glendenning and S. A. Moszkowski. Reconciliation of neutron-star masses and binding of the Λ in hypernuclei. *Physical Review Letters*, 67:2414–2417, October 1991.

[102] N. K. Glendenning and J. Schaffner-Bielich. First order kaon condensate. *Physical Review C*, 60:025803, August 1999.

[103] T. Gold. Rotating neutron stars as the origin of the pulsating radio sources. *Nature*, 218:731–732, May 1968.

[104] P. Grandclément, M. Ihm, V. Kalogera, and K. Belczynski. Searching for gravitational waves from the inspiral of precessing binary systems: Astrophysical expectations and detection efficiency of "spiky" templates. *Physical Review D*, 69:102002, May 2004.

[105] P. Grandclément and V. Kalogera. Searching for gravitational waves from the inspiral of precessing binary systems: New hierarchical scheme using "spiky" templates. *Physical Review D*, 67:082002, April 2003.

[106] P. Grandclément, V. Kalogera, and A. Vecchio. Searching for gravitational waves from the inspiral of precessing binary systems: Problems with current waveforms. *Physical Review D*, 67:042003, February 2003.

[107] J. B. Hartle. Slowly rotating relativistic stars. I. Equations of structure. *The Astrophysical Journal*, 150:1005–1030, December 1967.

[108] J. B. Hartle. Tidal friction in slowly rotating black holes. *Physical Review D*, 8:1010–1024, August 1973.

[109] J. B. Hartle. *Gravity: An Introduction to Einstein's General Relativity*. Addison Wesley, 2003.

[110] J. B. Hartle and K. S. Thorne. Slowly rotating relativistic stars. II. Models for neutron stars and supermassive stars. *The Astrophysical Journal*, 153:807–834, September 1968.

[111] S. W. Hawking. Gravitational radiation from colliding black holes. *Physical Review Letters*, 26:1344–1346, May 1971.

[112] A. Heger, C. L. Fryer, S. E. Woosley, N. Langer, and D. H. Hartmann. How massive single stars end their life. *The Astrophysical Journal*, 591:288–300, July 2003.

[113] R. W. Hellings and T. A. Moore. The information content of gravitational wave harmonics in compact binary inspiral. *Classical and Quantum Gravity*, 20:S181–S192, May 2003.

[114] J. W. T. Hessels, S. M. Ransom, I. H. Stairs, P. C. C. Freire, V. M. Kaspi, and F. Camilo. A radio pulsar spinning at 716 Hz. *Science*, 311:1901–1904, March 2006.

[115] A. Hewish, S. J. Bell, J. D. Pilkington, P. F. Scott, and R. A. Collins. Observation of a rapidly pulsating radio source. *Nature*, 217:709–713, February 1968.

[116] T. Hinderer. Tidal Love numbers of neutron stars. *The Astrophysical Journal*, 677:1216–1220, April 2008.

[117] T. Hinderer. Tidal Love numbers of neutron stars. arXiv:0711.2420v4, March 2009.

[118] G. Hinshaw, J. L. Weiland, R. S. Hill, N. Odegard, D. Larson, C. L. Bennett, J. Dunkley, B. Gold, M. R. Greason, N. Jarosik, E. Komatsu, M. R. Nolta, L. Page, D. N. Spergel, E. Wollack, M. Halpern, A. Kogut, M. Limon, S. S. Meyer, G. S. Tucker, and E. L. Wright. Five-year Wilkinson Microwave Anisotropy Probe observations: Data processing, sky maps, and basic results. *The Astrophysical Journal Supplement Series*, 180:225–245, February 2009.

[119] W. A. Hiscock. On tidal interactions with Kerr black holes. *The Astrophysical Journal*, 216:908–913, September 1977.

[120] D. W. Hogg. Distance measures in cosmology. arXiv:astro-ph/9905116, May 1999.

[121] D. E. Holz. Lensing and high-z supernova surveys. *The Astrophysical Journal*, 506:L1–L5, October 1998.

[122] D. E. Holz. Private communication, December 2008.

[123] D. E. Holz and S. A. Hughes. Using gravitational-wave standard sirens. *The Astrophysical Journal*, 629:15–22, August 2005.

[124] D. E. Holz and E. V. Linder. Safety in numbers: Gravitational lensing degradation of the luminosity distance-redshift relation. *The Astrophysical Journal*, 631:678–688, October 2005.

[125] D. E. Holz and R. M. Wald. New method for determining cumulative gravitational lensing effects in inhomogeneous universes. *Physical Review D*, 58:063501, September 1998.

[126] F. Hoyle and W. A. Fowler. Nucleosynthesis in supernovae. *The Astrophysical Journal*, 132:565–590, November 1960.

[127] S. A. Hughes. Untangling the merger history of massive black holes with LISA. *Monthly Notices of the Royal Astronomical Society*, 331:805–816, April 2002.

[128] S. A. Hughes. Listening to the universe with gravitational-wave astronomy. *Annals of Physics*, 303:142–178, January 2003.

[129] S. A. Hughes. (Sort of) Testing relativity with extreme mass ratio inspirals. In S. M. Merkovitz and J. C. Livas, editors, *Laser Interferometer Space Antenna: 6th International LISA Symposium*, volume 873 of *American Institute of Physics Conference Series*, pages 233–240, November 2006.

[130] S. A. Hughes. Gravitational waves from merging compact binaries. arXiv:0903.4877, March 2009.

[131] S. A. Hughes and R. D. Blandford. Black hole mass and spin coevolution by mergers. *The Astrophysical Journal*, 585:L101–L104, March 2003.

[132] S. A. Hughes and K. Menou. Golden binary gravitational-wave sources: Robust probes of strong-field gravity. *The Astrophysical Journal*, 623:689–699, April 2005.

[133] R. A. Hulse and J. H. Taylor. Discovery of a pulsar in a binary system. *The Astrophysical Journal*, 195:L51–L53, January 1975.

[134] I. Iben, Jr. and A. V. Tutukov. Supernovae of Type I as end products of the evolution of binaries with components of moderate initial mass ($M \leq 9M_{\odot}$). *The Astrophysical Journal Supplement Series*, 54:335–372, February 1984.

[135] A. H. Jaffe and D. C. Backer. Gravitational waves probe the coalescence rate of massive black hole binaries. *The Astrophysical Journal*, 583:616–631, February 2003.

[136] J. H. Jeans. The motion of tidally-distorted masses, with special reference to theories of cosmogony. *Memoirs of the Royal Astronomical Society*, 62:1–48, 1917.

[137] M. Kamionkowski, A. Kosowsky, and A. Stebbins. A probe of primordial gravity waves and vorticity. *Physical Review Letters*, 78:2058–2061, March 1997.

[138] K. Kawabe and the LIGO Collaboration. Status of LIGO. *Journal of Physics Conference Series*, 120:032003, July 2008.

[139] L. E. Kidder. Coalescing binary systems of compact objects to $(\text{post})^{5/2}$-Newtonian order. V. Spin effects. *Physical Review D*, 52:821–847, July 1995.

[140] L. E. Kidder, C. M. Will, and A. G. Wiseman. Spin effects in the inspiral of coalescing compact binaries. *Physical Review D*, 47:R4183–R4187, May 1993.

[141] C. S. Kochanek. Coalescing binary neutron stars. *The Astrophysical Journal*, 398:234–247, October 1992.

[142] B. Kocsis, Z. Frei, Z. Haiman, and K. Menou. Finding the electromagnetic counterparts of cosmological standard sirens. *The Astrophysical Journal*, 637:27–37, January 2006.

[143] B. Kocsis, Z. Haiman, and K. Menou. Premerger localization of gravitational wave standard sirens with LISA: Triggered search for an electromagnetic counterpart. *The Astrophysical Journal*, 684:870–887, September 2008.

[144] B. Kocsis, Z. Haiman, K. Menou, and Z. Frei. Premerger localization of gravitational-wave standard sirens with LISA: Harmonic mode decomposition. *Physical Review D*, 76:022003, July 2007.

[145] B. Kocsis and A. Loeb. Brightening of an accretion disk due to viscous dissipation of gravitational waves during the coalescence of supermassive black holes. *Physical Review Letters*, 101:041101, July 2008.

[146] K. Kokkotas and B. Schmidt. Quasi-normal modes of stars and black holes. *Living Reviews in Relativity*, 2:2, September 1999.

[147] S. Komossa, V. Burwitz, G. Hasinger, P. Predehl, J. S. Kaastra, and Y. Ikebe. Discovery of a binary active galactic nucleus in the ultraluminous infrared galaxy NGC 6240 using Chandra. *The Astrophysical Journal*, 582:L15–L19, January 2003.

[148] J. Kormendy and D. Richstone. Inward bound: The search for supermassive black holes in galactic nuclei. *Annual Review of Astronomy and Astrophysics*, 33:581–624, September 1995.

[149] S. Kuroyanagi, T. Chiba, and N. Sugiyama. Precision calculations of the gravitational wave background spectrum from inflation. *Physical Review D*, 79:103501, May 2009.

[150] I. E. Lagaris and V. R. Pandharipande. Variational calculations of realistic models of nuclear matter. *Nuclear Physics A*, 359:349–364, April 1981.

[151] D. Lai. Tidal stablization of neutron stars and white dwarfs. *Physical Review Letters*, 76:4878–4881, June 1996.

[152] D. Lai, F. A. Rasio, and S. L. Shapiro. Ellipsoidal figures of equilibrium: Compressible models. *The Astrophysical Journal Supplement Series*, 88:205–252, September 1993.

[153] D. Lai, F. A. Rasio, and S. L. Shapiro. Hydrodynamics of rotating stars and close binary interactions: Compressible ellipsoid models. *The Astrophysical Journal*, 437:742–769, December 1994.

[154] R. N. Lang and S. A. Hughes. Measuring coalescing massive binary black holes with gravitational waves: The impact of spin-induced precession. *Physical Review D*, 74:122001, December 2006.

[155] R. N. Lang and S. A. Hughes. Erratum: Measuring coalescing massive binary black holes with gravitational waves: The impact of spin-induced precession [Phys. Rev. D 74, 122001 (2006)]. *Physical Review D*, 75:089902, April 2007.

[156] R. N. Lang and S. A. Hughes. Erratum: Measuring coalescing massive binary black holes with gravitational waves: The impact of spin-induced precession [Phys. Rev. D 74, 122001 (2006)]. *Physical Review D*, 77:109901, May 2008.

[157] R. N. Lang and S. A. Hughes. Localizing coalescing massive black hole binaries with gravitational waves. *The Astrophysical Journal*, 677:1184–1200, April 2008.

[158] R. N. Lang and S. A. Hughes. Advanced localization of massive black hole coalescences with LISA. *Classical and Quantum Gravity*, 26:094035, May 2009.

[159] S. L. Larson, W. A. Hiscock, and R. W. Hellings. Sensitivity curves for spaceborne gravitational wave interferometers. *Physical Review D*, 62:062001, September 2000.

[160] J. M. Lattimer. Private communication, April 2009.

[161] J. M. Lattimer and M. Prakash. Neutron star structure and the equation of state. *The Astrophysical Journal*, 550:426–442, March 2001.

[162] J. M. Lattimer and M. Prakash. The physics of neutron stars. *Science*, 304:536–542, April 2004.

[163] J. M. Lattimer and M. Prakash. Neutron star observations: Prognosis for equation of state constraints. *Physics Reports*, 442:109–165, April 2007.

[164] U. J. Le Verrier. Théorie du mouvement de Mercure. *Annales de l'Observatoire Impérial de Paris*, 5:1–195, 1859.

[165] Z. Lippai, Z. Frei, and Z. Haiman. Prompt shocks in the gas disk around a recoiling supermassive black hole binary. *The Astrophysical Journal*, 676:L5–L8, March 2008.

[166] A. G. Lyne, M. Burgay, M. Kramer, A. Possenti, R. N. Manchester, F. Camilo, M. A. McLaughlin, D. R. Lorimer, N. D'Amico, B. C. Joshi, J. Reynolds, and P. C. C. Freire. A double-pulsar system: A rare laboratory for relativistic gravity and plasma physics. *Science*, 303:1153–1157, February 2004.

[167] A. I. MacFadyen and M. Milosavljević. An eccentric circumbinary accretion disk and the detection of binary massive black holes. *The Astrophysical Journal*, 672:83–93, January 2008.

[168] J. Magorrian, S. Tremaine, D. Richstone, R. Bender, G. Bower, A. Dressler, S. M. Faber, K. Gebhardt, R. Green, C. Grillmair, J. Kormendy, and T. Lauer. The demography of massive dark objects in galaxy centers. *The Astronomical Journal*, 115:2285–2305, June 1998.

[169] F. K. Manasse and C. W. Misner. Fermi normal coordinates and some basic concepts in differential geometry. *Journal of Mathematical Physics*, 4:735–745, June 1963.

[170] J.-A. Marck. Solution to the equations of parallel transport in Kerr geometry; tidal tensor. *Proceedings of the Royal Society A*, 385:431–438, 1983.

[171] J. M. Martín-García and C. Gundlach. Gauge-invariant and coordinate-independent perturbations of stellar collapse. II. Matching to the exterior. *Physical Review D*, 64:024012, July 2001.

[172] B. Mashhoon. On tidal phenomena in a strong gravitational field. *The Astrophysical Journal*, 197:705–716, May 1975.

[173] N. Mavalvala. Private communication, May 2009.

[174] D. L. Meier. The theory and simulation of relativistic jet formation: towards a unified model for micro- and macroquasars. *New Astronomy Review*, 47:667–672, October 2003.

[175] M. Micic, K. Holley-Bockelmann, S. Sigurdsson, and T. Abel. Supermassive black hole growth and merger rates from cosmological N-body simulations. *Monthly Notices of the Royal Astronomical Society*, 380:1533–1540, October 2007.

[176] M. Milosavljević and D. Merritt. The final parsec problem. In J. M. Centrella, editor, *The Astrophysics of Gravitational Wave Sources*, volume 686 of *American Institute of Physics Conference Series*, pages 201–210, October 2003.

[177] M. Milosavljević and E. S. Phinney. The afterglow of massive black hole coalescence. *The Astrophysical Journal*, 622:L93–L96, April 2005.

[178] T. Mora and C. M. Will. Post-Newtonian diagnostic of quasiequilibrium binary configurations of compact objects. *Physical Review D*, 69:104021, May 2004.

[179] H. Müller and B. D. Serot. Relativistic mean-field theory and the high-density nuclear equation of state. *Nuclear Physics A*, 606:508–537, September 1996.

[180] H. Müther, M. Prakash, and T. L. Ainsworth. The nuclear symmetry energy in relativistic Brueckner-Hartree-Fock calculations. *Physics Letters B*, 199:469–474, December 1987.

[181] R. Narayan, B. Paczynski, and T. Piran. Gamma-ray bursts as the death throes of massive binary stars. *The Astrophysical Journal*, 395:L83–L86, August 1992.

[182] J. W. Negele and D. Vautherin. Neutron star matter at sub-nuclear densities. *Nuclear Physics A*, 207:298–320, June 1973.

[183] G. Nelemans, L. R. Yungelson, and S. F. Portegies Zwart. The gravitational wave signal from the Galactic disk population of binaries containing two compact objects. *Astronomy & Astrophysics*, 375:890–898, September 2001.

[184] J. R. Oppenheimer and G. M. Volkoff. On massive neutron cores. *Physical Review*, 55:374–381, February 1939.

[185] S. O'Sullivan. Private communication, April 2009.

[186] Y. Pan, A. Buonanno, J. G. Baker, J. Centrella, B. J. Kelly, S. T. McWilliams, F. Pretorius, and J. R. van Meter. Data-analysis driven comparison of analytic and numerical coalescing binary waveforms: Nonspinning case. *Physical Review D*, 77:024014, January 2008.

[187] Y. Pan, A. Buonanno, Y. Chen, and M. Vallisneri. Physical template family for gravitational waves from precessing binaries of spinning compact objects: Application to single-spin binaries. *Physical Review D*, 69:104017, May 2004.

[188] V. R. Pandharipande and R. A. Smith. A model neutron solid with π^0 condensate. *Nuclear Physics A*, 237:507–532, January 1975.

[189] S. Perlmutter, G. Aldering, G. Goldhaber, R. A. Knop, P. Nugent, P. G. Castro, S. Deustua, S. Fabbro, A. Goobar, D. E. Groom, I. M. Hook, A. G. Kim, M. Y. Kim, J. C. Lee, N. J. Nunes, R. Pain, C. R. Pennypacker, R. Quimby, C. Lidman, R. S. Ellis, M. Irwin, R. G. McMahon, P. Ruiz-Lapuente, N. Walton, B. Schaefer, B. J. Boyle, A. V. Filippenko, T. Matheson, A. S. Fruchter, N. Panagia, H. J. M. Newberg, W. J. Couch, and The Supernova Cosmology Project. Measurements of Ω and Λ from 42 high-redshift supernovae. *The Astrophysical Journal*, 517:565–586, June 1999.

[190] M. E. Peskin and D. V. Schroeder. *An Introduction to Quantum Field Theory*. Westview Press, 1995.

[191] P. C. Peters. Gravitational radiation and the motion of two point masses. *Physical Review*, 136:1224–1232, November 1964.

[192] M. M. Phillips. The absolute magnitudes of Type Ia supernovae. *The Astrophysical Journal*, 413:L105–L108, August 1993.

[193] E. Poisson. Private communication, July 2008.

[194] E. Poisson and C. M. Will. Gravitational waves from inspiraling compact binaries: Parameter estimation using second-post-Newtonian waveforms. *Physical Review D*, 52:848–855, July 1995.

[195] S. F. Portegies Zwart and S. L. W. McMillan. The runaway growth of intermediate-mass black holes in dense star clusters. *The Astrophysical Journal*, 576:899–907, September 2002.

[196] E. K. Porter and N. J. Cornish. Effect of higher harmonic corrections on the detection of massive black hole binaries with LISA. *Physical Review D*, 78:064005, September 2008.

[197] R. A. Porto and I. Z. Rothstein. Calculation of the first nonlinear contribution to the general-relativistic spin-spin interaction for binary systems. *Physical Review Letters*, 97:021101, July 2006.

[198] M. Prakash, J. R. Cooke, and J. M. Lattimer. Quark-hadron phase transition in protoneutron stars. *Physical Review D*, 52:661–665, July 1995.

[199] M. Prakash, J. M. Lattimer, and T. L. Ainsworth. Equation of state and the maximum mass of neutron stars. *Physical Review Letters*, 61:2518–2521, November 1988.

[200] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, 1992.

[201] F. Pretorius. Evolution of binary black-hole spacetimes. *Physical Review Letters*, 95:121101, September 2005.

[202] R. Price and K. S. Thorne. Non-radial pulsation of general-relativistic stellar models. II. Properties of the gravitational waves. *The Astrophysical Journal*, 155:163–182, January 1969.

[203] E. Rantsiou, S. Kobayashi, P. Laguna, and F. A. Rasio. Mergers of black hole-neutron star binaries. I. Methods and first results. *The Astrophysical Journal*, 680:1326–1349, June 2008.

[204] T. Regge and J. A. Wheeler. Stability of a Schwarzschild singularity. *Physical Review*, 108:1063–1069, November 1957.

[205] C. S. Reynolds and M. A. Nowak. Fluorescent iron lines as a probe of astrophysical black hole systems. *Physics Reports*, 377:389–466, April 2003.

[206] A. G. Riess, A. V. Filippenko, P. Challis, A. Clocchiatti, A. Diercks, P. M. Garnavich, R. L. Gilliland, C. J. Hogan, S. Jha, R. P. Kirshner, B. Leibundgut, M. M. Phillips, D. Reiss, B. P. Schmidt, R. A. Schommer, R. C. Smith, J. Spyromilio, C. Stubbs, N. B. Suntzeff, and J. Tonry. Observational evidence from supernovae for an accelerating universe and a cosmological constant. *The Astronomical Journal*, 116:1009–1038, September 1998.

[207] A. G. Riess, W. H. Press, and R. P. Kirshner. Using Type Ia supernova light curve shapes to measure the Hubble constant. *The Astrophysical Journal*, 438:L17–L20, January 1995.

[208] É. Roche. Mémoire sur la figure d'une masse fluide (soumise à l'attraction d'un point éloigné). *Académie des Sciences et Lettres de Montpellier*, 1:243–262, 333–348, 1847-50.

[209] C. Rodriguez, G. B. Taylor, R. T. Zavala, A. B. Peck, L. K. Pollack, and R. W. Romani. A compact supermassive binary black hole system. *The Astrophysical Journal*, 646:49–60, July 2006.

[210] L. J. Rubbo. *Gravitational Wave Astronomy Using Spaceborne Detectors*. PhD thesis, Montana State University, July 2004.

[211] A. Sadowski, K. Belczynski, T. Bulik, N. Ivanova, F. A. Rasio, and R. O'Shaughnessy. The total merger rate of compact object binaries in the local universe. *The Astrophysical Journal*, 676:1162–1169, April 2008.

[212] M. V. Sazhin. Opportunities for detecting ultralong gravitational waves. *Soviet Astronomy*, 22:36–38, February 1978.

[213] H. R. Schmitt, J. E. Pringle, C. J. Clarke, and A. L. Kinney. The orientation of jets relative to dust disks in radio galaxies. *The Astrophysical Journal*, 575:150–155, August 2002.

[214] J. D. Schnittman. Spin-orbit resonance and the evolution of compact binary systems. *Physical Review D*, 70:124020, December 2004.

[215] J. D. Schnittman, A. Buonanno, J. R. van Meter, J. G. Baker, W. D. Boggs, J. Centrella, B. J. Kelly, and S. T. McWilliams. Anatomy of the binary black hole recoil: A multipolar analysis. *Physical Review D*, 77:044031, February 2008.

[216] J. D. Schnittman and J. H. Krolik. The infrared afterglow of supermassive black hole mergers. *The Astrophysical Journal*, 684:835–844, September 2008.

[217] B. F. Schutz. Determining the Hubble constant from gravitational wave observations. *Nature*, 323:310–311, September 1986.

[218] U. Seljak and M. Zaldarriaga. Signature of gravity waves in the polarization of the microwave background. *Physical Review Letters*, 78:2054–2057, March 1997.

[219] A. Sesana, F. Haardt, P. Madau, and M. Volonteri. The gravitational wave signal from massive black hole binaries and its contribution to the LISA data stream. *The Astrophysical Journal*, 623:23–30, April 2005.

[220] A. Sesana, M. Volonteri, and F. Haardt. The imprint of massive black hole formation models on the LISA data stream. *Monthly Notices of the Royal Astronomical Society*, 377:1711–1716, June 2007.

[221] S. L. Shapiro and S. A. Teukolsky. *Black Holes, White Dwarfs, and Neutron Stars: The Physics of Compact Objects*. Wiley-Interscience, 1983.

[222] M. Shibata. Relativistic Roche-Riemann problems around a black hole. *Progress of Theoretical Physics*, 96:917–932, November 1996.

[223] M. Shibata, K. Kyutoku, T. Yamamoto, and K. Taniguchi. Gravitational waves from black hole-neutron star binaries: Classification of waveforms. *Physical Review D*, 79:044030, February 2009.

[224] M. Shibata and K. Taniguchi. Merger of black hole and neutron star in general relativity: Tidal disruption, torus mass, and gravitational waves. *Physical Review D*, 77:084015, April 2008.

[225] M. Shibata and K. Uryū. Merger of black hole-neutron star binaries: Nonspinning black hole case. *Physical Review D*, 74:121503, December 2006.

[226] M. Shibata and K. Uryū. Merger of black hole neutron star binaries in full general relativity. *Classical and Quantum Gravity*, 24:S125–S137, June 2007.

[227] G. A. Shields and E. W. Bonning. Powerful flares from recoiling black holes in quasars. *The Astrophysical Journal*, 682:758–766, August 2008.

[228] S. Sigurdsson and M. J. Rees. Capture of stellar mass compact objects by massive black holes in galactic cusps. *Monthly Notices of the Royal Astronomical Society*, 284:318–326, January 1997.

[229] D. N. Spergel, R. Bean, O. Doré, M. R. Nolta, C. L. Bennett, J. Dunkley, G. Hinshaw, N. Jarosik, E. Komatsu, L. Page, H. V. Peiris, L. Verde, M. Halpern, R. S. Hill, A. Kogut, M. Limon, S. S. Meyer, N. Odegard, G. S. Tucker, J. L. Weiland, E. Wollack, and E. L. Wright. Three-year Wilkinson Microwave Anisotropy Probe (WMAP) observations: Implications for cosmology. *The Astrophysical Journal Supplement Series*, 170:377–408, June 2007.

[230] V. Springel, S. D. M. White, A. Jenkins, C. S. Frenk, N. Yoshida, L. Gao, J. Navarro, R. Thacker, D. Croton, J. Helly, J. A. Peacock, S. Cole, P. Thomas, H. Couchman, A. Evrard, J. Colberg, and F. Pearce. Simulations of the formation, evolution and clustering of galaxies and quasars. *Nature*, 435:629–636, June 2005.

[231] R. T. Stebbins. Rightsizing LISA. *Classical and Quantum Gravity*, 26:094014, May 2009.

[232] K. Taniguchi, T. W. Baumgarte, J. A. Faber, and S. L. Shapiro. Relativistic black hole-neutron star binaries in quasiequilibrium: Effects of the black hole excision boundary condition. *Physical Review D*, 77:044003, February 2008.

[233] T. M. Tauris and E. van den Heuvel. Formation and evolution of compact stellar X-ray sources. arXiv:astro-ph/0303456, March 2003.

[234] J. H. Taylor and J. M. Weisberg. A new test of general relativity: Gravitational radiation and the binary pulsar PSR 1913+16. *The Astrophysical Journal*, 253:908–920, February 1982.

[235] J. H. Taylor and J. M. Weisberg. Further experimental tests of relativistic gravity using the binary pulsar PSR 1913+16. *The Astrophysical Journal*, 345:434–450, October 1989.

[236] S. Taylor and E. Poisson. Nonrotating black hole in a post-Newtonian tidal environment. *Physical Review D*, 78:084016, October 2008.

[237] K. S. Thorne. Multipole expansions of gravitational radiation. *Reviews of Modern Physics*, 52:299–340, April 1980.

[238] K. S. Thorne. Tidal stabilization of rigidly rotating, fully relativistic neutron stars. *Physical Review D*, 58:124031, December 1998.

[239] K. S. Thorne and A. Campolattaro. Non-radial pulsation of general-relativistic stellar models. I. Analytic analysis for $\ell \geq 2$. *The Astrophysical Journal*, 149:591–612, September 1967.

[240] K. S. Thorne and J. B. Hartle. Laws of motion and precession for black holes and other bodies. *Physical Review D*, 31:1815–1837, April 1985.

[241] K. S. Thorne, R. H. Price, and D. A. MacDonald. *Black Holes: The Membrane Paradigm*. Yale University Press, 1986.

[242] J. I. Thorpe, S. T. McWilliams, B. J. Kelly, R. P. Fahey, K. Arnaud, and J. G. Baker. LISA parameter estimation using numerical merger waveforms. *Classical and Quantum Gravity*, 26:094026, May 2009.

[243] B. G. Todd-Rutel and J. Piekarewicz. Neutron-rich nuclei and neutron stars: A new accurately calibrated interaction for the study of neutron-rich matter. *Physical Review Letters*, 95:122501, September 2005.

[244] R. C. Tolman. Static solutions of Einstein's field equations for spheres of fluid. *Physical Review*, 55:364–373, February 1939.

[245] M. Trias and A. M. Sintes. LISA observations of supermassive black holes: Parameter estimation using full post-Newtonian inspiral waveforms. *Physical Review D*, 77:024030, January 2008.

[246] J. A. Tyson. Large Synoptic Survey Telescope: Overview. In J. A. Tyson and S. Wolff, editors, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 4836, pages 10–20, December 2002.

[247] R. Umstätter, R. Meyer, R. J. Dupuis, J. Veitch, G. Woan, and N. Christensen. Estimating the parameters of gravitational waves from neutron stars using an adaptive MCMC method. *Classical and Quantum Gravity*, 21:S1655–S1665, October 2004.

[248] M. Vallisneri. Prospects for gravitational-wave observations of neutron-star tidal disruption in neutron-star–black-hole binaries. *Physical Review Letters*, 84:3519–3522, April 2000.

[249] M. J. Valtonen. New orbit solutions for the precessing binary black hole model of OJ 287. *The Astrophysical Journal*, 659:1074–1081, April 2007.

[250] A. Vecchio. LISA observations of rapidly spinning massive black hole binary systems. *Physical Review D*, 70:042001, August 2004.

[251] L. Wang, G. Goldhaber, G. Aldering, and S. Perlmutter. Multicolor light curves of Type Ia supernovae on the color-magnitude diagram: A novel step toward more precise distance and extinction estimates. *The Astrophysical Journal*, 590:944–970, June 2003.

[252] R. F. Webbink. Double white dwarfs as progenitors of R Coronae Borealis stars and Type I supernovae. *The Astrophysical Journal*, 277:355–360, February 1984.

[253] J. Whelan and I. J. Iben. Binaries and supernovae of Type I. *The Astrophysical Journal*, 186:1007–1014, December 1973.

[254] P. Wiggins and D. Lai. Tidal interaction between a fluid star and a Kerr black hole in circular orbit. *The Astrophysical Journal*, 532:530–539, March 2000.

[255] C. M. Will and A. G. Wiseman. Gravitational radiation from compact binary systems: Gravitational waveforms and energy loss to second post-Newtonian order. *Physical Review D*, 54:4813–4848, October 1996.

[256] R. E. Williams, B. Blacker, M. Dickinson, W. V. D. Dixon, H. C. Ferguson, A. S. Fruchter, M. Giavalisco, R. L. Gilliland, I. Heyer, R. Katsanis, Z. Levay, R. A. Lucas, D. B. McElroy, L. Petro, M. Postman, H.-M. Adorf, and R. Hook. The Hubble Deep Field: Observations, data reduction, and galaxy photometry. *The Astronomical Journal*, 112:1335–1384, October 1996.

[257] B. Willke, P. Aufmuth, C. Aulbert, S. Babak, R. Balasubramanian, B. W. Barr, S. Berukoff, S. Bose, G. Cagnoli, M. M. Casey, D. Churches, D. Clubley, C. N. Colacino, D. R. M. Crooks, C. Cutler, K. Danzmann, R. Davies, R. Dupuis, E. Elliffe, C. Fallnich, A. Freise, S. Goßler, A. Grant, H. Grote, G. Heinzel, A. Heptonstall, M. Heurs, M. Hewitson, J. Hough, O. Jennrich, K. Kawabe, K. Kötter, V. Leonhardt, H. Lück, M. Malec, P. W. McNamara, S. A. McIntosh, K. Mossavi, S. Mohanty, S. Mukherjee, S. Nagano, G. P. Newton, B. J. Owen, D. Palmer, M. A. Papa, M. V. Plissi, V. Quetschke, D. I. Robertson, N. A. Robertson, S. Rowan, A. Rüdiger, B. S. Sathyaprakash, R. Schilling, B. F. Schutz, R. Senior, A. M. Sintes, K. D. Skeldon, P. Sneddon, F. Stief, K. A. Strain, I. Taylor, C. I. Torrie, A. Vecchio, H. Ward, U. Weiland, H. Welling, P. Williams, W. Winkler, G. Woan, and I. Zawischa. The GEO 600 gravitational wave detector. *Classical and Quantum Gravity*, 19:1377–1387, April 2002.

[258] R. B. Wiringa, V. Fiks, and A. Fabrocini. Equation of state for dense nucleon matter. *Physical Review C*, 38:1010–1037, August 1988.

[259] E. Witten. Cosmic separation of phases. *Physical Review D*, 30:272–285, July 1984.

[260] T. Yamamoto, M. Shibata, and K. Taniguchi. Simulating coalescing compact binaries by a new code (SACRA). *Physical Review D*, 78:064054, September 2008.

[261] N. Yunes, W. Tichy, B. J. Owen, and B. Brügmann. Binary black hole initial data from matched asymptotic expansions. *Physical Review D*, 74:104011, November 2006.

[262] X.-H. Zhang. Multipole expansions of the general-relativistic gravitational field of the external universe. *Physical Review D*, 34:991–1004, August 1986.