

“The nitrogen in our DNA, the calcium in our teeth, the iron in our blood, the carbon in our apple pies were made in the interiors of collapsing stars.

We are made of star stuff.”

Carl Sagan

Cosmos

University of Alberta

Understanding Parameter Degeneracies in Neutron Star X-ray Light Curves

by

Abigail Lauren Stevens

A thesis submitted to the Faculty of Graduate Studies and Research
in partial fulfillment of the requirements for the degree of

Master of Science

Department of Physics

© Abigail Lauren Stevens
Fall 2013
Edmonton, Alberta

Permission is hereby granted to the University of Alberta Libraries to reproduce single copies of this thesis and to lend or sell such copies for private, scholarly or scientific research purposes only. Where the thesis is converted to, or otherwise made available in digital form, the University of Alberta will advise potential users of the thesis of these terms.

The author reserves all other publication and other rights in association with the copyright in the thesis and, except as herein before provided, neither the thesis nor any substantial portion thereof may be printed or otherwise reproduced in any material form whatsoever without the author's prior written permission.

Abstract

The equation of state for cold ultra-dense matter has puzzled astrophysicists for decades. This is because the conditions of supra-nuclear density matter, such as those in neutron stars, are not terrestrially replicable. X-ray light curves from accreting neutron stars have proven to be useful tools in studying the neutron star equation of state. Theory predicts that the light curve from a thermonuclear X-ray burst on a rapidly-rotating neutron star can be used to determine the characteristics of the burst ignition spot and constrain the mass and radius of the neutron star. We discuss the development of spherical and oblate neutron star models that yield an X-ray light curve comparable to that which would be observed. Using this simulation code with a genetic algorithm, we disentangle the effects of various parameters on the light curve, showing which parameter degeneracies will have the greatest impact on the observables.

Acknowledgements

This thesis research would not have been possible without the assistance and encouragement of those who have helped me along the way. **Thank you** to **Sharon Morsink** for introducing me to the wonderful world of neutron star physics and for your guidance and support as my supervisor; to **Jason Fiege** for collaboration on setting up and implementing the Ferret genetic algorithm; to **Denis Leahy** for collaboration on producing the simulated data and for providing independent results for comparison; to my committee members **John Davis**, **Craig Heinke**, and **Gregory Sivakoff** for advice and feedback on the development of this thesis; to **Adam Rogers** for invaluable feedback on chapter 4; to **Faun Rice** for assistance organizing and editing chapter 3; to **Robin Arnason**, **Arash Bahramian**, **Jeanette Gladstone**, **Sarah Nowicki**, **Logan Sibley**, and **Tania Wood** for numerous fruitful conversations contributing to the development of my research and this thesis; to the **Department of Physics** and the **astrophysics research group** for helping me to grow as an academic, as a scientist, and as a person; to my **family** for your love and support; to my **friends** and **Tyler** for wonderful adventures, advice, and awesomeness, and for making Edmonton feel like home for the past two years; and to my **housemates** at the Golden Lentils for being wonderful to come home to and for enriching my life. Finally, thank you to the **Faculty of Graduate Studies and Research**, the **Graduate Student Association**, and the **CASCA** Board of Directors for funding opportunities that allowed me to travel to conferences to present this research, and to **NSERC** for funding my studies and research in Summer 2012 and Summer 2013.

Contents

1	Introduction	1
1.1	Neutron Stars	1
1.1.1	Equation of State	2
1.2	X-ray Binaries	3
1.2.1	Accretion in LMXBs	4
1.2.2	Accreting Millisecond X-ray Pulsars	5
1.3	Type I X-ray Bursts	6
1.3.1	Thermonuclear Burst Oscillations	7
1.4	Light Curves	8
1.4.1	History of the Method	8
1.5	X-ray Timing Telescopes	16
1.5.1	RXTE	16
1.5.2	ASTROSAT	17
1.5.3	NICER	17
1.5.4	LOFT	17
1.6	Organization of this Thesis	18
2	Mathematical Theory	19
2.1	Metric and Approximation	20
2.2	Light Bending	22
2.3	Time of Arrival	27
2.4	Doppler Effects	27

2.5	Blackbody Radiation	30
2.6	Oblate Approximation	31
3	Light Curve Simulation Code	32
3.1	Hotspot Parameters	33
3.1.1	Essential Parameters	33
3.1.2	Additional Parameters	36
3.2	Non-Trivially-Sized Hotspots	38
3.3	Flux in Energy Bands	41
3.4	Hotspot Shapes	45
3.5	Non-Zero Background Temperature	51
3.6	Two Antipodal Hotspots	51
3.6.1	Low Inclination	52
3.6.2	Moderate Inclination	53
3.6.3	High Inclination	54
3.7	Hotspot Covering the Rotation Axis	54
3.7.1	Symmetric Over the Rotation Axis	55
3.7.2	Asymmetric Over the Rotation Axis	56
4	Fitting with a Genetic Algorithm	57
4.1	The Concept of Genetic Algorithms	58
4.2	Qubist and Ferret	61
4.2.1	Components of Ferret	65
4.3	Other Fitting Methods	67
4.3.1	MCMC + Bayesian	67
4.3.2	Levenberg-Marquardt	69
5	Understanding Parameter Degeneracies	70
5.1	Characterizing Light Curves	71
5.2	Degeneracy Contour Code	73

5.2.1	Solving for Inclination and Emission Angles	73
5.2.2	Components of the Degeneracy Contour Code	76
5.3	Parameter Degeneracies in Our Research	76
6	Results	80
6.1	Case 1: The Fiducial Case	81
6.1.1	Ferret	82
6.1.2	Parameter Degeneracy	89
6.2	Case 2: Strong Doppler Boost	95
6.2.1	Ferret	95
6.2.2	Parameter Degeneracy	97
6.3	Case 3: High Noise, Symmetric, Very Low Amplitude	97
6.3.1	Ferret	97
6.3.2	Parameter Degeneracy	97
6.4	Case 4: Symmetric, Moderate Amplitude	98
6.4.1	Ferret	99
6.4.2	Parameter Degeneracy	100
6.5	Case 5: Graybody Model	101
6.5.1	Ferret	102
6.5.2	Parameter Degeneracy	102
6.6	Case 6: Oblate Model	104
6.6.1	Ferret	104
6.6.2	Parameter Degeneracy	106
6.7	Case 7: Oblate, Graybody Model	106
6.7.1	Ferret	106
6.7.2	Parameter Degeneracy	106
6.8	Case 8: Moderate Doppler Boost	107
6.8.1	Ferret	108
6.8.2	Parameter Degeneracy	109

6.9	Case 9: Symmetric, Moderately Low Amplitude	110
6.9.1	Ferret	111
6.9.2	Parameter Degeneracy	112
6.10	Case 10: Fiducial, High Noise	113
6.10.1	Ferret	114
6.10.2	Parameter Degeneracy	115
6.11	Discussion and Summary	115
7	Conclusion	121
7.1	Future Research	122
	Appendix	124

List of Tables

- 5.1 Pulse amplitudes 73
- 6.1 Results, Case 1 85
- 6.3 Results, all 94
- 6.4 Phase shift correlation with σ deviation 116
- A.1 Sample 15×15 temperature mesh 124

List of Figures

2.1	Inclination angle	22
2.2	Emission angle and hotspot angular size	23
2.3	Phase angle	24
2.4	Geometry of angles	25
2.5	Impact parameter	26
3.1	Varying inclination angle	34
3.2	Varying emission angle	35
3.3	Sample 4×4 mesh over a hotspot	40
3.4	Mesh over small hotspot	41
3.5	Mesh over large hotspot	42
3.6	Telescope energy bands	43
3.7	Normal energy bands	44
3.8	Planck function of 2 keV hotspot	46
3.9	Doppler boost factor and hotspot solid angle vs. phase	47
3.10	High energy bands	48
3.11	Different hotspot shapes	49
3.12	Hotspot shapes	50
3.13	Two hotspots, low inclination	53
3.14	Two hotspots, moderate inclination	54
3.15	Two hotspots, high inclination	55
5.1	Light curve characteristics	71

5.2	Example of parameter degeneracy	78
6.1	Case 1: Light curves, low energy band	83
6.2	Case 1: Light curves, high energy band	84
6.3	Case 1: Ferret χ^2 contours	86
6.4	Case 1: χ^2 surface, 3D	87
6.5	Case 1: χ^2 surface, 2D colour map	88
6.6	Case 1: χ^2 parameter degeneracy contours	90
6.7	Case 1: Overlapping χ^2 contour regions	91
6.8	Case 2: Light curves	95
6.9	Case 2: χ^2 contours	96
6.10	Case 3: Light curves	98
6.11	Case 3: χ^2 contours	99
6.12	Case 4: Light curves	100
6.13	Case 4: χ^2 contours	101
6.14	Case 5: Light curves	102
6.15	Case 5: χ^2 contours	103
6.16	Case 6: Light curves	104
6.17	Case 6: χ^2 contours	105
6.18	Case 7: Light curves	107
6.19	Case 7: χ^2 contours	108
6.20	Case 8: Light curves	109
6.21	Case 8: χ^2 contours	110
6.22	Case 9: Light curves	111
6.23	Case 9: χ^2 contours	112
6.24	Case 10: Light curves	113
6.25	Case 10: χ^2 contours	114
6.26	Graybody and oblateness	119

Chapter 1

Introduction

1.1 Neutron Stars

Neutron stars are the ultracompact remnants of moderately massive stars that were initially $\sim 8 - 20$ times the mass of our sun (M_{\odot}). First theorized in 1933 by Baade and Zwicky [1], neutron stars were discovered by Hewish, Bell, et al. in 1967 [2]. Neutron stars form in core-collapse supernovae. They have masses on the order of $1 M_{\odot}$, radii on the order of 10 km, internal temperatures < 1 MeV a few seconds after birth, and strong magnetic fields ranging from $\sim 10^7$ to 10^{15} G [3][4][5]. The acceleration due to gravity on the surface of a neutron star is $\sim 10^{12} \text{ m s}^{-2}$ [5].

Initial calculations of theoretical neutron star models by Oppenheimer and Volkoff in 1939 assumed a highly compressed ideal gas of free neutrons [6]. However, while neutron stars have approximately no net electric charge [1], they cannot solely be a dense sea of pure neutrons, since this configuration is very unstable. Modern observations and theories allow us to better predict their physical composition. The outer crust of a neutron star is theoretically composed of an iron Coulomb lattice [1], where highly ionized iron atoms share electrons in a lattice structure and Coulomb forces are assumed to dominate the interactions [5]. The inner crust is composed of densely packed neutron-

rich elements like heavy iron, heavy nickel, and heavy krypton [7]. As we move radially inward, densities increase to greater than neutron drip density, $\sim 4 \times 10^{11} \text{ g cm}^{-3}$, at which point it is energetically favourable for the neutrons to leave their home nucleus and they “drip” off of the atoms [7]. Much of the core of a neutron star is thought to consist of superfluid neutrons [8][9], with minor amounts of superconducting protons and relativistic electrons [7]. The inner core is not well-understood, since it has a density greater than an atomic nucleus, $\sim 10^{15} \text{ g cm}^{-3}$ (nuclear density is $\sim 2.8 \times 10^{14} \text{ g cm}^{-3}$) [1][5]. There are many theories about what this material is (e.g., quark matter, hyperons, various particle condensates), but this active area of research in neutron star physics has not yet yielded definitive answers [1].

As a first approximation, a neutron star is supported against gravitational collapse by neutron degeneracy pressure, with further corrections given by the strong nuclear force. The degeneracy pressure is a direct consequence of the Heisenberg uncertainty principle. The compactness of a neutron star confines the position of the neutrons, which causes the momentum of the neutrons to have a larger uncertainty; this yields a pressure between the particles much greater than the thermal pressure [7]. The equilibrium of pressure and gravity that keeps the neutron star physically stable is referred to as hydrostatic equilibrium.

1.1.1 Equation of State

The holy grail in neutron star research is to determine the equation of state for neutron star matter. An equation of state is a way to quantitatively describe the behaviour of matter in a given physical state. We commonly write equations of state as a relation between pressure, density, and temperature, but since the thermal energy of neutrons in a neutron star is below the Fermi energy, $\sim 30 \text{ MeV}$, we can neglect the temperature [1]. Equations of state are known for many types of matter, but not for cold supranuclear-density matter. For

neutron star matter, different equations of state will allow for different sizes of neutron stars. These so-called “soft” and “stiff” equations of state are relative. Soft equations of state have a lower pressure at a given density, and tend to give smaller neutron stars with a smaller maximum mass. Stiff equations of state have a higher pressure at a given density, and tend to allow for larger and more massive neutron stars. Additionally, rapid rotation ($\gtrsim 300$ Hz) can allow for a more massive neutron star at a given radius, since the centrifugal effects help to further support the neutron star structure against gravitational collapse [4][10].

Obtaining constraints on the masses and radii of neutron stars narrows the possible equations of state to describe ultra-dense nuclear matter in the inner cores of neutron stars. We use neutron stars to learn about the ultra-dense cold equation of state since this physics is not terrestrially replicable.

1.2 X-ray Binaries

A binary system consists of two stars gravitationally bound to each other in orbit around a common center of mass. Binaries are considered X-ray binaries when they are bright in the X-ray regime, due to the release of energy from matter falling down in the gravitational potential well of a compact object (accretion). Of particular interest are X-ray binaries with a neutron star accretor. We classify these binaries in terms of the mass of the companion as high mass X-ray binaries (HMXBs) or low mass X-ray binaries (LMXBs) [3]. In HMXBs, the companion is an O- or B-type star. Matter is transferred from the companion to the primary typically through stellar wind mass loss from the companion. In LMXBs, the companion is typically less massive than the Sun. Here, matter is transferred by the companion overflowing its Roche Lobe (volume of space in which the companion is gravitationally dominant), at which point the overflowing matter falls towards the accretor [7]. Due to the angular momentum of

the infalling matter, the accretor accumulates matter from the companion in a disk about the accretor’s rotational equator. As accumulated matter orbits in the accretion disk, it moves in closer to the accretor. Since material at smaller disk radii makes more orbits per unit time than than material farther out in the disk, there can be viscous heating between different rings of the accretion disk. This heating causes the emission of X-ray radiation [3]. Therefore the presence of X-ray emission informs us that the binary is undergoing active accretion. For our purposes, emission from the accretion disk is background X-ray emission coming from the system.

1.2.1 Accretion in Low Mass X-ray Binaries

In a LMXB, the neutron star accretes mass from the companion via Roche Lobe overflow. Stable mass transfer happens onto the accretion disk on nuclear timescales (the time a nuclear supply lasts, $\sim 10^{10} \frac{M}{M_{\odot}} \frac{L_{\odot}}{L}$ years) [11], however accretion onto the neutron star can occur in a stable or unstable regime. We assume a simple Ghosh and Lamb accretion disk model, in which the disk extends inward until approximately the co-rotation point [12], where the orbiting material has the same rotational frequency as the neutron star’s spin (and magnetic field). Here, material in the innermost part of the accretion disk moves along the magnetic field lines onto the surface of the neutron star. As the neutron star accumulates matter from the accretion disk, it gains the angular momentum of this matter. Over time, this can spin up the neutron star to millisecond spin periods [13][14]. These are referred to as “recycled” neutron stars. We expect to see LMXBs in globular clusters such as 47 Tucanae, where the stellar interaction rate is very high [15].

1.2.2 Accreting Millisecond X-ray Pulsars

A pulsar is a rotating neutron star that typically emits in the radio regime. The regular pulsed emission pattern is caused by observing the emitting region as it faces towards and away from the observer during a rotation. Two types of pulsars are known: radio pulsars, in which the emission is powered by rotation, and X-ray pulsars, which are powered by accretion.

Accreting millisecond X-ray pulsars are rapidly spinning accretion-powered pulsars with spin periods of a few milliseconds. Their pulsed X-ray emission originates from material slamming onto the surface of the neutron star during regular accretion and warming the surface so that it emits blackbody radiation. Then, as the pulsar rotates, it gives periodic oscillations in brightness as the emitting region faces towards and away from the observer. The first known accreting millisecond X-ray pulsar, SAX J1808.4–3658 (hereafter, SAX J1808), was observed by Wijnands and van der Klis in 1998 [13]. Accreting millisecond X-ray pulsars are thought to be the missing link between LMXBs and millisecond radio pulsars. Papitto et al. 2013 [16] have recently observed a known millisecond radio pulsar switching off and becoming an accreting millisecond X-ray pulsar, suggesting that rapidly rotating pulsars can switch between these two states on a timescale of a few years. It is believed that a millisecond pulsar will alternate between radio and X-ray states until its companion stops transferring matter, at which point it becomes exclusively a millisecond radio pulsar. This new evidence strengthens the claim that accreting millisecond X-ray pulsars are a missing link between LMXBs and millisecond radio pulsars.

Moreover, accreting millisecond X-ray pulsars are thought to have stronger magnetic fields than regular neutron stars in LMXBs, $\sim 10^{10} - 10^{12}$ G [17]. In theory, this leads to accreted material funnelling onto the magnetic poles of the neutron star, and strong magnetic fields over the poles could suspend matter in an accretion column above the surface of the neutron star.

In addition to the regular X-ray emission, some of the phenomena these

objects have are bigger, better, and more explosive: Type I X-ray bursts.

1.3 Type I X-ray Bursts

The accreted material building up on the surface of the neutron star stratifies based on density. Once it reaches a critical density and temperature it ignites unstable thermonuclear burning, and releases a large burst of X-rays [18]. These localized thermonuclear X-ray bursts are categorized as Type I X-ray bursts. Typically we expect to see helium and possibly hydrogen signatures in the burst spectrum [19], since the neutron star is usually accreting from a main sequence star or a white dwarf. Observed Type I X-ray bursts have peak luminosities $\sim 10^{38} \text{ erg s}^{-1}$ and last a few seconds to minutes [17]. These X-ray bursts are significantly brighter than the “background” X-ray emission from the accretion disk.

The ignition and emission area on the surface of the neutron star is referred to as the hotspot. It is presently unknown if material is only accreting at the hotspot for these neutron stars with relatively low magnetic fields, or if the hotspot is the only area where thermonuclear ignition and X-ray emission occurs. Theories suggest two different surface hotspot models: one that ignites nuclear burning at one point and spreads across the whole star, and another that ignites at one point and begins to spread but remains limited to a smaller area (referred to as a “persistent hotspot”) [18]. We will show models for the second type of hotspot once they have established the persistent burning region, but these models could be applied to one snapshot in time of the spreading hotspot model, as long as the neutron star’s spin period is much less than the time it takes for the burning to spread significantly. The persistent hotspot on the surface of a rotating neutron star has been demonstrated to be an effective model for Type I X-ray bursts from the source 4U 1636–536 [20].

The neutron star’s magnetic field might get buried in the accumulation of

matter on the surface in systems that exhibit Type I X-ray bursts, which would explain why the neutron stars in these systems have low inferred magnetic fields ($\sim 10^8$ – 10^9 G) compared to other neutron stars [21]. If magnetic fields are strong enough, such as in accreting millisecond X-ray pulsars, accreted material will accumulate at the magnetic poles, which likely would give two antipodal hotspots, one at the north magnetic pole and another at the south magnetic pole. While some accreting millisecond X-ray pulsars also exhibit Type I X-ray bursts, not all Type I X-ray bursts come from accreting millisecond X-ray pulsars.

1.3.1 Thermonuclear Burst Oscillations

In some Type I X-ray bursts we observe brightness oscillations whose frequency strongly corresponds with the spin frequency of the neutron star; these are referred to as thermonuclear burst oscillations. Thermonuclear burst oscillations are not detected in all Type I X-ray bursts, nor are they detected in every Type I X-ray burst from sources that have exhibited thermonuclear burst oscillations in the past. Additionally, thermonuclear burst oscillations are not detected during an entire Type I X-ray burst. Typically, burst oscillations are seen in the rise and/or tail of the X-ray burst, and not during the peak (e.g. see figures in Muno et al. 2002 [22] and Watts 2012 [18]). The amplitudes of thermonuclear burst oscillations vary greatly during the X-ray burst, with high amplitude oscillations found in the rise of the burst and lower amplitude oscillations found in the tail [22][23]. Furthermore, only thermonuclear burst oscillations from Type I X-ray bursts on accreting millisecond X-ray pulsars have shown some harmonic content such as asymmetric pulse shapes; burst oscillations from X-ray bursts on neutron stars without persistent pulsations tend to have negligible harmonic contributions to the oscillation pulse shape [22].

Most burst oscillations seen in the rise of the X-ray burst exhibit frequency

drift, where the frequency of the oscillations increases by 1–3 Hz towards the known spin frequency of the neutron star [18][22]. Frequency drift is a common complication in modelling burst oscillations in the rise of an X-ray burst, but this thesis assumes a simplified burst oscillation model which does not address these concerns.

There are presently 17 sources that exhibit thermonuclear burst oscillations (see Watts 2012 [18] and references therein).

1.4 Light Curves

One of the strongest theories to explain thermonuclear burst oscillations is the hotspot model, since the burst oscillation frequency approximately matches the known frequency of the neutron star [18][23]. We expect that the hotspot is not aligned with the rotation axis of the neutron star. This misalignment gives way to the “lighthouse effect”, so that as the neutron star rotates, we see varying amounts of photons from the X-ray burst at the hotspot. This gives rise to observed periodic X-ray emission, or thermonuclear burst oscillations. Our research focuses on simulating X-ray emission from hotspots in thermonuclear burst oscillations, which this thesis will refer to as a light curve.

Since these X-rays originate from the surface of the neutron star, the light curve contains information about the physical properties of the neutron star itself. Modelling a light curve is referred to as “ray-tracing” the photon from the surface of the neutron star out to an observer at infinity. By modelling light curves and fitting them to observed data, we hope to constrain the masses and radii of neutron stars.

1.4.1 History of the Method

In this section we review previous implementations of light curve ray-tracing and fitting (of thermonuclear burst oscillations from Type I X-ray bursts and of

persistent pulsations from accreting millisecond X-ray pulsars) in the literature.

The canonical paper by Pechenick et al. 1983 [24] is the first to report the formalism for computing photon trajectories from a neutron star surface embedded in a Schwarzschild spacetime (typically used to describe the curved spacetime around a non-rotating black hole). This method accounts for gravitational light bending and lensing of the strong gravitational field surrounding the neutron star, but neglects frame dragging terms from a rotating spacetime and rotational effects such as Doppler boosting. They assume that the emission is only from the hotspot, and that there are no perturbations to the photons once they are emitted from the neutron star’s surface. Since this includes assuming no magnetic field effects on the photon paths, this method is most directly applicable to neutron stars with relatively weak magnetic fields, $\sim 10^7$ – 10^{12} G [25]. Initially this method was to be used strictly for isotropic thermal emission, but they include a rudimentary beaming pattern, where emission is not isotropically emitted.

The two hotspot geometries shown by Pechenick et al. are for one polar cap or two polar caps, and the computations for the bending angles we show in Chapter 2 were introduced to the literature in this paper. The authors explore various parameter geometries, showing their resultant light curves for different emission set-ups. One of the key points demonstrated by this paper is that eclipses in the light curve, in which the hotspot is not visible to the observer even when accounting for light bending and so the observed flux is zero, are not commonly expected. Additionally, the authors note that due to gravitational light bending, “no point more than 2 or 3 km above the surface of a typical neutron star is ever out of sight” [24]. This has important implications in column models, where the emitting region is a column of material suspended above the surface of the neutron star. However, the light curves that have been observed from accreting millisecond X-ray pulsars are fairly simple and featureless, and so they do not require accretion columns to get good fits [26][27].

The photon trajectory ray-tracing mathematical approach is the basis or inspiration for nearly all light curve methods seen in the literature. It has been adapted to model a variety of emitting objects like anomalous X-ray pulsars, neutron stars in HMXBs, emission of light near stellar mass black holes, or observed data; object properties like two non-antipodal spots, various atmospheres, global oscillation modes, or complicated hotspot shapes; and emission mechanisms like inverse Compton scattering.

Understanding emission from isolated pulsars has benefitted from the Pechenick et al. light curve formalism. Harding and Muslimov 1998 [28] model X-ray light curves of the known isolated gamma ray pulsars Geminga, Vela, and PSR 0656+14 using different beaming models of thermal emission along the magnetic field lines. DeDeo et al. [29] re-create observed light curve shapes and features seen in data from anomalous X-ray pulsar observations. As a final example, Perna and Gotthelf 2008 [30] re-create anomalous X-ray pulsar data with a double blackbody spectrum, appropriate for a warm outer ring and hot inner cap emission area.

Emission from neutron stars in HMXBs is also a viable application of this light curve modelling, typically using accretion columns instead of hotspot caps. Kraus 2001 [31] apply a hollow accretion column model to Centaurus X-3. Leahy 2004 [32] fits observed light curves from Hercules X-1 using an accretion column model, with a fan-shaped beam emission pattern close to the surface of the neutron star and a pencil-shaped beam emission pattern further from the surface. Sasaki et al. 2010 [33] also model beam patterns to fit light curves, and find that a single emission region is the best fit for EXO 2030+375. Annala and Poutanen 2010 [34] find that using a pencil-shaped beam emission pattern, the majority of slowly rotating X-ray pulsars in HMXBs likely have accretion onto a small spot at the magnetic poles. As a final example, Caballero et al. 2011 [35] fit observations of A 0535+26 with two non-antipodal hot spots that each have a hollow accretion column and a surrounding scattering “halo”

region.

Furthermore, the light curve computation method outlined by Pechenick et al. sets the foundation for modelling emission from outside the event horizon of black holes. Psaltis and Johannsen 2012 [36] adapt the method to use a Kerr metric for stellar mass black holes, and Chan et al. 2013 [37] introduce GRay, a massive parallel GPU-based integrator that traces the trajectories of photons around Kerr black holes.

The Pechenick et al. ray-tracing method has been a competent model of observed data of persistent X-ray emission from accreting millisecond X-ray pulsars. Nath et al. 2002 [38] fit models of two known accreting millisecond X-ray pulsars with Type I X-ray bursts, 4U 1636–536 and 4U 1728–34, using a Schwarzschild metric and electron scattering atmosphere but neglecting Doppler effects. Morsink and Leahy 2011 [27] fit data from SAX J1808 from six different epochs of emission, under the assumption that the mass and radius of the neutron star are constant from one epoch to the next. As a final example, Lo et al. 2013 [39] use a Schwarzschild metric with Doppler effects (we refer to this as the “Schwarzschild+Doppler” method) and test out a fitting approach using a Markov chain Monte Carlo algorithm (discussed in more detail in Section 4.3.1).

We can also use this method to explore what the light curve would look like from two non-antipodal hotspots, as if from a neutron star with off-center magnetic poles. Weinberg et al. 2001 [40] model emission from two non-antipodal hotspots resulting in asymmetric light curves. Bogdanov et al. 2007 [41] (revisited in Bogdanov 2013 [42]) explain the thermal emission from PSR J043–4715 with an off-center magnetic dipole causing two non-antipodal emitting regions.

The effects of various atmospheres on the shape of the light curve is fairly well-documented in the literature as a continuation on the Pechenick et al. formalism. Meszaros and Riffert 1988 [43] model cap and column emission models with photon propagation via radiative transfer in an atmosphere and then

standard Schwarzschild spacetime photon propagation to the observer. Muno et al. 2003 [44] fit light curves from Type I X-ray bursts from 4U 1636–536 with a model assuming that photons emitted from the surface are scattered by a hot corona of electrons before reaching the observer, and find that this “graybody”-type atmosphere is more likely than pure thermal emission. Bhattacharyya et al. 2005 [45] fit asymmetric Type I X-ray burst light curves from XTE J1848–338 with a Kerr metric and Doppler effects, and find that the photons emitted at the surface experience Thomson scattering in the accreted layers on top of the emission site as they travel to the observer. Ho 2007 [46] explores the effects a magnetic hydrogen atmosphere would have on emission from two antipodal hot spots at the magnetic poles. As a final example, Gotthelf et al. 2010 [47] model double blackbody emission from two antipodal hotspots on Puppis A and place constraints on the magnetic field strength and system geometry.

Global oscillation modes, small physical oscillations of spherical harmonics of the neutron star matter, are thought to possibly play a role in neutron star X-ray burst emission patterns. Strohmayer 1992 [48] models the anisotropic beaming from large polar caps on a slowly rotating neutron star with global oscillation modes. Lee and Strohmayer 2005 [49] calculate theoretical light curves produced by different oscillation modes propagating in a fluid ocean along the equator of a neutron star (not from hotspots), using a Schwarzschild+Doppler method. As a final example, Numata and Lee 2010 [50] compute light curves as would be produced from a hotspot on a rapidly rotating neutron star if the hotspot is perturbed by certain oscillation mode waves.

Methods of photon propagation have also been coupled to other codes to model complicated hotspot shapes and how the resulting light curve would appear. Kulkarni and Romanova 2005 [51] use a pseudo-Newtonian+Doppler photon propagation method in combination with 3-dimensional magneto-hydrodynamic simulations of accretion onto a neutron star surface, and predict real-

istic hotspot shapes.

Another emission effect typically unaccounted for is inverse-Compton scattering, in which photons scatter off of energetic electrons to higher frequencies [52]. Viironen and Poutanen 2004 [53] model emission from two antipodal hotspots from accreting millisecond X-ray pulsars and Type I X-ray bursts, looking at how inverse-Compton scattering off a hot electron cloud above the hotspot would affect the spectral profile of initially thermal emission.

We now discuss some key adaptations to the Pechenick et al. methodology which set the stage for our work.

The basis of our particular technique first appeared in Miller and Lamb 1998 [54]. This uses the Schwarzschild photon propagation seen in [24], and also accounts for the effects of gravitational redshift and Doppler boosting of rapid rotation on the observed light curve. Their sample models have a neutron star spin frequency of 600 Hz, for which the rotational Doppler effects are nontrivial. As in [24], the authors use a point-like hotspot cap geometry to model light curves for one or two hotspots. The motivation for developing the Schwarzschild+Doppler model is to analyze brightness oscillations in the tails of Type I X-ray bursts, which can be seen in the light curves from rapidly rotating neutron stars. The goal of the authors in this paper is to determine the largest possible amplitude of oscillation for a given neutron star compactness.

Our method is most directly based off Poutanen and Gierlinski 2003 [55]. They use the Schwarzschild+Doppler model from [54], but replace the full ray-tracing computation with the Beloborodov approximation for calculating the photon's trajectory (shown in equation (2.2.6)) [56]. Additionally, they expand the method's applicability to include a non-trivially-sized hotspot and an inverse-Compton scattering region above the hotspot on the surface. The authors fit observed light curves from SAX J1808 to constrain the parameters of the neutron star and emission geometry. Their fitting procedure is not explicitly defined, but they are the first paper that uses a method similar to ours to fit a

model to an observed light curve to constrain neutron star parameters. From their best fit models the authors are able to constrain the neutron star radius and observer inclination angle of SAX J1808.

Cadeau et al. 2005 [57] explain the importance of computing photon paths with the full ray-tracing integral versus the Beloborodov approximation, and of accounting for time delays between photon emission and detection for rapidly rotating neutron stars. This is the first paper in the literature that outlines the necessity of accounting for time delay effects. The authors compare the output of light curve models with and without time delay, concluding that neglecting time delays can moderately increase the errors for inferred parameters. Time delay is explained in further detail in Section 2.3.

Cadeau et al. 2007 [58] outline the need for an oblate neutron star shape, particularly for rapidly rotating neutron stars (~ 600 Hz), using the formalism from [24][54][55]. The oblate neutron star model includes centrifugal effects from rapid rotation, which yields a neutron star with a larger radius at its equator than along the spin axis. Note that oblateness is directly caused by a high spin frequency. All previous models discussed assume a spherical neutron star shape, but the authors show that this is not as realistic an approximation when compared to numerically solving the correct metric in equation (2.1.1). Furthermore, they demonstrate that the difference between a spherical neutron star in a Kerr metric (which accounts for frame dragging) and spherical neutron star in a Schwarzschild metric is not nearly as drastic as the difference between spherical neutron star models and oblate neutron star models. This paper discusses the geometry of the oblate formalism and how one could constrain the radius of a neutron star based on the amount of Doppler boosting seen in the light curve. The authors conclude that a Schwarzschild+Doppler model with an oblate neutron star shape is a more reasonable approximation for rapidly rotating neutron stars than the spherical shape.

The full exposition of the oblate Schwarzschild+Doppler formalism appears

in Morsink et al. 2007 [59]. This work is motivated by Cadeau et al. 2007 [58], based on Pechenick et al. 1983 [24], Miller and Lamb 1998 [54], and Poutanen and Gierlinski 2003 [55]. Here the authors lay out the mathematical details for photon emission from a rapidly rotating oblate spheroid embedded in Schwarzschild spacetime, accounting for time delay effects. Using point-like spots, they have two oblate models: one for neutron and hybrid quark stars, and another for colour-flavour locked quark stars. In addition, the authors show how the effects of oblateness can be conflated with Doppler effects for some system geometries.

Using the spherical Schwarzschild+Doppler approach, Lamb et al. 2009 [60] re-create the very low pulse amplitudes seen in accretion powered emission from accreting millisecond X-ray pulsars. The authors discuss a model for a hotspot that moves in latitude and longitude on the emission timescale, as well a hotspot that moves on longer time scales. The emission regions modelled here are very close to the rotational axes. It is thought that the spin up of the neutron star to ~ 400 Hz causes the magnetic poles to move close to the rotational axes [60]. While not explicitly fitting to data, the authors discuss how different parameters (hotspot size, emission angle, and inclination angle) can affect the light curve, showing that many different combinations of parameters can give the desired low pulse amplitude. This is one of the first extensive discussions of parameter degeneracies in the literature that uses models and methods very similar to ours.

The most recent application of the oblate “Schwarzschild+Doppler” formalism is in Leahy et al. 2011 [26]. The authors accommodate inverse-Compton emission as well as the standard blackbody emission for the observed flux in different energy bands from an accreting millisecond X-ray pulsar, noting that inverse-Compton scattering only affects higher energy bands. With this, they are able to fit the observed time-independent spectrum as well as the periodic light curve. Their model for point-like hot spots has an anisotropy parameter

to allow beam shapes from inverse-Compton scattering and an additional photon scattering region between the observer and the neutron star. The authors use this model to constrain the parameters of XTE J1807–294, fitting with the Levenberg-Marquardt algorithm (discussed in detail in Section 4.3.2).

1.5 X-ray Timing Telescopes

Since we are looking at phenomena that are happening on millisecond timescales, and since Type-I X-ray bursts emit in the X-ray regime (as the name suggests), we require a space-based X-ray timing telescope with sub-millisecond timing resolution for observations. Four relevant satellites are the former NASA mission RXTE, the future Indian Space Research Organization mission ASTROSAT, the future NASA mission NICER, and the proposed ESA mission LOFT.

1.5.1 Rossi X-ray Timing Explorer

The Rossi X-ray Timing Explorer, RXTE, is a former NASA X-ray timing satellite. It was launched into low-Earth circular orbit on December 30, 1995, and was decommissioned on January 5, 2012 [61]. The Proportional Counter Array (PCA), comprised of five proportional photon counters, was the most applicable instrument to our research. It had a time resolution of $1\ \mu\text{s}$, an effective area of $0.65\ \text{m}^2$ at 6 keV, and a spectral resolution of 1 keV. The key sensitivity range was 2–25 keV, but extended up to 60 keV [62].

Much of what is known about Type I X-ray bursts and other short-timescale phenomena on compact objects comes from observations made with RXTE [25][63][64][65].

1.5.2 ASTROSAT

ASTROSAT is an Indian Space Research Organisation (ISRO) multi-wavelength observing mission with three instruments for X-ray observing. Additionally, it will feature a wide-field all sky monitor and two optical-to-UV-band telescopes [66]. Currently the projected launch date is aimed for 2014.

The Large Area Xenon Proportional Counter will be best suited to our research. It will have a timing resolution of $10 \mu\text{s}$, an effective area of 0.6 m^2 , a spectral resolution of 9% at 22 keV, and a sensitivity range of 3–80 keV [66]. Note that ASTROSAT will have a higher sensitivity than RXTE at energies $> 15 \text{ keV}$.

1.5.3 Neutron star Interior Composition ExploreR

NICER, the Neutron star Interior Composition ExploreR, is a recently approved NASA X-ray telescope slated for launch in 2016 [67]. It will be attached to the International Space Station. The instrument will observe in the 0.2–12 keV X-ray band with a time resolution of $0.3 \mu\text{s}$, an effective area of 0.06 m^2 at 10 keV, and a spectral resolution of 0.137 keV [67]. This small, specialized instrument will allow for further observations of short-timescale X-ray phenomena when it is launched in late 2016 [67].

1.5.4 Large Observatory For x-ray Timing

LOFT, the Large Observatory For x-ray Timing, is a proposed European Space Agency (ESA) X-ray timing satellite mission as a part of their Cosmic Vision program, addressing the topic of “matter under extreme conditions” [68]. This medium-sized mission candidate is currently in the assessment phase along with four other candidates, and in a few years one of them will be chosen to launch in about 2022. LOFT would be a next-generation X-ray timing mission to expand on the science that was made possible by RXTE.

The two scientific components of LOFT are the Large Area Detector (LAD) and the Wide-Field Monitor (WFM). The LAD is a collimated (pointable) instrument with an expected effective area of 12 m^2 in the 2–10 keV energy range. Its projected time resolution is $10 \mu\text{s}$, its spectral resolution is 0.26 keV, and its key sensitivity range is 2–30 keV with sensitivity extending to 50 keV [68]. The WFM is intended for source monitoring and detection of interesting intensity and spectral states for various burst sources. It has a 3.95 sr field of view, 6 arcmin angular resolution, and 2 mCrab sensitivity (at 5σ after 50 ks) [68].

The two main science goals of the LOFT team are to determine if matter orbiting close to the event horizon of a black hole follows the predictions of general relativity, and to constrain the equation of state of ultradense matter in neutron stars. The research in this thesis is contributing to the dense matter group in placing constraints on the masses and radii of neutron stars.

1.6 Organization of this Thesis

Chapter 2 covers the mathematical theory involved in the research, such as the spacetime metric used, gravitational light bending effects, the time of arrival computation, Doppler effects, blackbody radiation, and the oblate neutron star approximation. Chapter 3 discusses the light curve simulation code and our additions to it. Chapter 4 discusses what a genetic algorithm is and how we fit light curves with the Ferret genetic algorithm. Chapter 5 discusses the parameter degeneracies we encounter and how to account for them. Chapter 6 shows the results from fitting light curves with a genetic algorithm and how these lie in the parameter space when accounting for parameter degeneracies. Chapter 7 is the conclusion and discusses future research possibilities. Appendix A has a sample temperature mesh over a large hotspot.

Chapter 2

Mathematical Theory

In this chapter we discuss the mathematical and physical theory behind the observed phenomena. The first part of the chapter covers the complete spacetime metric and the Schwarzschild approximation. The second part of the chapter includes equations relating to the gravitational light bending of photon emission from the surface of a neutron star. The third part of the chapter covers the time of arrival of emitted photons reaching an observer at infinity. The fourth part of the chapter contains equations relating to the Doppler boosting of emission due to the rapid rotation of the neutron star. The fifth part of the chapter covers equations relating to the blackbody emission from the neutron star, as well as the pulse amplitude measured by an observer at infinity. The sixth part of the chapter discusses the adaptations required for an oblate neutron star.

All radii are measured from the center of the neutron star to the surface at the hotspot unless otherwise stated. The subscript “*em*” means the parameter is in the neutron star’s rest frame, the subscript “*s*” means the parameter is for a static neutron star measured by an observer at infinity, and the subscript “*obs*” means the parameter accounts for Doppler boosting as observed by an observer at infinity. The equations in this chapter are from Poutanen and Gierlinski 2003 [55] and Morsink et al. 2007 [59]. We only include the key

equations.

Most values in this chapter use gravitational units, where $G = c = 1$. For example, to transform mass and radius from SI to gravitational units, we use

$$\frac{2GM}{Rc^2} = \frac{2M_{gr}}{R_{gr}}, \quad (2.0.1)$$

however we commonly drop the subscript “ gr ” when writing equations, even when using gravitational units. The gravitational unit for angles is radians. If a value is intended to have standard units, it will be explicitly stated as such.

2.1 Metric and Approximation

The stationary axisymmetric metric for a rotating neutron star is [58][69]

$$ds^2 = -e^{\gamma+\Xi} dt^2 + e^{2\alpha} (d\bar{r}^2 + \bar{r}^2 d\theta^2) + e^{\gamma-\Xi} \bar{r}^2 \sin^2 \theta (d\phi - \omega dt)^2, \quad (2.1.1)$$

where γ , Ξ , α , and ω are metric potentials, or descriptors of the gravitational field. These depend on the coordinates \bar{r} and θ . The coordinates are defined as follows: ds is a spacetime interval, t is the time coordinate, \bar{r} is a modified radial coordinate, θ is the polar angular coordinate, and ϕ is the azimuthal angular coordinate. The isotropic Schwarzschild radius, \bar{r} , is used here since equation (2.1.1) does not limit to the Schwarzschild metric coordinate system for a non-rotating star. This \bar{r} and the typical radial coordinate r are related by the coordinate transform $r = \bar{r} [1 + M/(2\bar{r})]^2$ (for reference, $2\pi r$ is the circumference of circle, while $2\pi\bar{r}$ is not) [70].

This metric requires a numerical approximation [10]. As shown in Cadeau et al. 2007 [58] we can instead approximate this using the simpler Schwarzschild metric (typically used to model a static, spherically symmetric black hole) and

accounting for the rotational Doppler effects. The Schwarzschild metric is [70]

$$ds^2 = -\left(1 - \frac{2M}{r}\right) dt^2 + \frac{1}{\left(1 - \frac{2M}{r}\right)} dr^2 + r^2 d\theta^2 + r^2 \sin^2 \theta d\phi^2, \quad (2.1.2)$$

where M is the mass of the neutron star, and the coordinates are the same as in equation (2.1.1), but with the standard radial coordinate r . This metric is preferable to the first because it can be solved exactly instead of numerically approximated. However, the Schwarzschild metric has a coordinate singularity at $r = 2M$, which lies inside the neutron star. Traditionally this is the location of the event horizon of a black hole.

The spacetime inside a spherical star is given by the metric [70]

$$ds^2 = -e^{2\Phi} dt^2 + e^{2\Lambda} dr^2 + r^2 d\theta^2 + r^2 \sin^2 \theta d\phi^2, \quad (2.1.3)$$

where Φ and Λ are gravitational field potentials as functions of r . We embed the neutron star surface, with this metric inside, in a Schwarzschild spacetime. The gravitational field potential Φ , which in the Newtonian limit is the Newtonian gravitational potential [70], is defined by the relativistic equations of stellar structure, [4]

$$\frac{dm}{dr} = 4\pi r^2 \rho, \quad (2.1.4)$$

$$\frac{dP}{dr} = -\frac{\rho m}{r^2} \left(1 + \frac{P}{\rho}\right) \left(1 + \frac{4\pi P r^3}{m}\right) \left(1 - \frac{2m}{r}\right)^{-1}, \quad (2.1.5)$$

$$\frac{d\Phi}{dr} = -\frac{1}{\rho} \frac{dP}{dr} \left(1 + \frac{P}{\rho}\right)^{-1}, \quad (2.1.6)$$

where ρ is density, P is pressure, r is the traditional radial coordinate, and m is a function of r for the mass contained inside a spherical shell at r . These interior equations are matched at the surface of the neutron star with the Schwarzschild metric so that $r = R$, where R is the radius of the neutron star, and $m(r) = m(R) = M$, where M is the mass of the neutron star. The

Newtonian limit for equations (2.1.4), (2.1.5), and (2.1.6) takes $P \ll \rho$ and $m \ll r$.

In addition to approximating the exterior spacetime with a Schwarzschild metric, we also account for Doppler effects of gravitational redshift and the Lorentz boost from the rapidly rotating neutron star (discussed in section 2.4). Together this mathematical framework comprises the Schwarzschild + Doppler method [54][55].

2.2 Light Bending

The incredible density of a neutron star causes it to bend spacetime, requiring the metric in the previous section. Bending spacetime produces a curved geodesic, the shortest distance between two points, which in flat spacetime is a straight line. Photons travel along geodesics, and so the path of a photon is curved in the space time around neutron stars; this is referred to as “light bending”. We will discuss the geometry of this and the equations needed to account for this phenomenon.

The inclination angle, shown in Figure 2.1, is defined as the angle between the rotation axis of the neutron star and the observer’s line of sight.

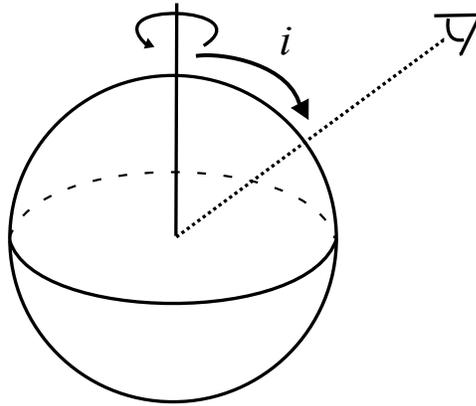


Figure 2.1: The inclination angle is the angle between the rotation axis of the neutron star and the observer’s line of sight.

The emission angle, θ , is measured at the center of the neutron star between the radial vector at the rotation axis and the radial vector going through the center of the emitting region, or “hotspot” (Figure 2.2). The angular radius of the hotspot, ρ , is measured at the center of the neutron star between the radial vector pointing to the center of the hotspot and the radial vector pointing to the outer edge of the hotspot.

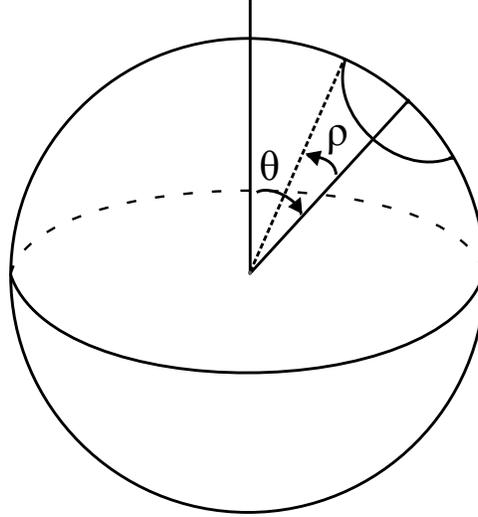


Figure 2.2: The emission angle, θ , is between the rotation axis of the neutron star and the center of the hotspot. The hotspot’s angular radius, ρ , is measured from the center of the hotspot to the outer edge of the hotspot.

The phase angle, ϕ , is measured in the equatorial around the star as it rotates, such that $\phi : 0 - 2\pi$ is one complete phase (Figure 2.3). Computationally, one rotation of the star has n phase bins, and the flux values in a light curve are computed once per phase bin.

Three key vectors (see Figure 2.4) are \mathbf{k} , which points in the direction of the observer located at infinity, \mathbf{l} (that is, lower-case “L”), which points in the initial photon direction, and \mathbf{r} , which points radially outward from where the photon was emitted. The zenith angle α is between the neutron star’s radial vector (the zenith of the hotspot) and direction in which photon is emitted:

$$\cos \alpha = \mathbf{l} \cdot \mathbf{r} . \tag{2.2.1}$$

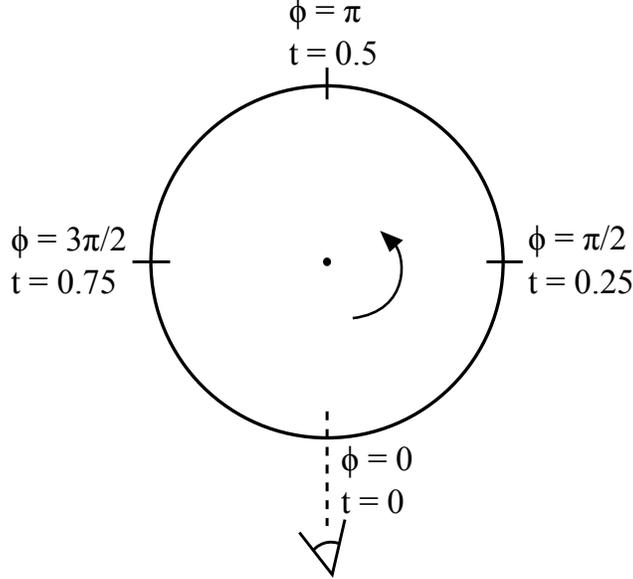


Figure 2.3: Top-down view of a neutron star, with the rotation axis coming out of the page at the center. The phase angle, ϕ , and normalized time, t , are noted for four positions. The dotted line indicates the observer's line of sight; we have defined ϕ such that the observer always sits at $\phi = 0$. The arrow indicates the direction in which the neutron star is rotating. We notice that at $t = 0.25$ the neutron star surface is moving radially away from the observer, and at $t = 0.75$ the neutron star surface is moving radially towards the observer. Since both phase angle and normalized time are cyclical parameters, $\phi = 0 = 2\pi$ and $t = 0 = 1$.

The bending angle, ψ , between the zenith of the hotspot and the observer's line of sight, is defined by spherical trigonometry to be [55][59]

$$\cos \psi = \mathbf{k} \cdot \mathbf{r} = \cos i \cos \theta + \sin i \sin \theta \cos \phi. \quad (2.2.2)$$

We only calculate photons that end up at the observer, so all photons calculated must move from \mathbf{l} to \mathbf{k} . In a flat spacetime, the photon is not able to bend, which means that only the photons emitted in the direction of the observer will reach the observer, and so $\psi = \alpha$ (as in Figure 2.5). Note that the vectors \mathbf{r} , \mathbf{k} and \mathbf{l} lie in the same plane, since the Schwarzschild metric requires a photon to move in one plane. Due to this, \mathbf{l} can be defined in terms of α , ψ , \mathbf{k} , and \mathbf{r}

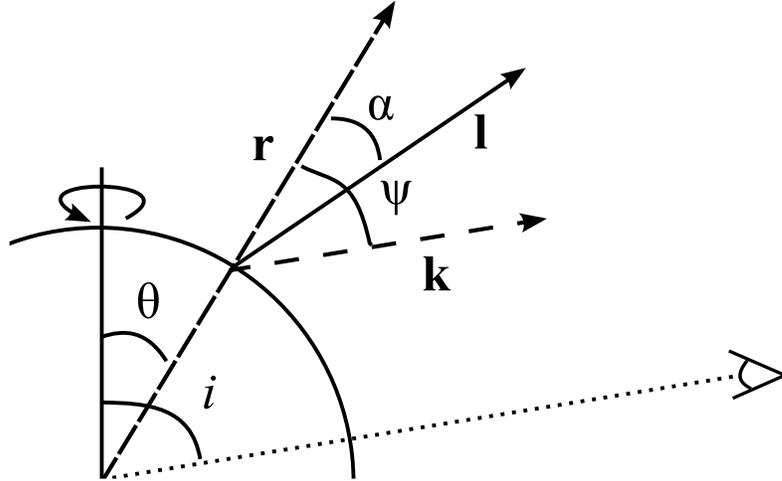


Figure 2.4: A close-up of a neutron star emitting a photon from a point-like hotspot on its surface. The vector \mathbf{r} points radially out of the neutron star at the hotspot. The vector \mathbf{l} points in the direction of the emission of the photon. The vector \mathbf{k} points in the direction of the observer located at infinity, so that this vector and the observer's line of sight (dotted line) are parallel. The zenith angle α is between the radial vector \mathbf{r} and the photon emission direction \mathbf{l} . The bending angle ψ is between the radial vector \mathbf{r} and the direction of the observer's line of sight \mathbf{k} .

such that [59]

$$\mathbf{l} = \frac{1}{\sin \psi} (\sin \alpha \mathbf{k} + \sin(\psi - \alpha) \mathbf{r}) . \quad (2.2.3)$$

These angles and vectors are shown together in Figure 2.4.

The bending angle ψ is computed using the full ray-tracing equation, integrating from the surface of the star to the observer at infinity, by [24]

$$\psi = b \int_R^\infty \frac{dr}{r^2} \left(1 - \frac{b^2}{r^2} \left(1 - \frac{2M}{r} \right) \right)^{-1/2} . \quad (2.2.4)$$

This equation comes from solving the geodesic equations for photon orbits in a Schwarzschild metric. The variable b is known as the impact parameter of a photon (Figure 2.5).

As implied by Figure 2.5, in Newtonian gravity we would expect $\sin \alpha = b/R$. In general relativity, the gravitational light bending effectively makes the neutron star appear to have a larger radius $R' = R/\sqrt{1 - 2M/R}$. This gives the

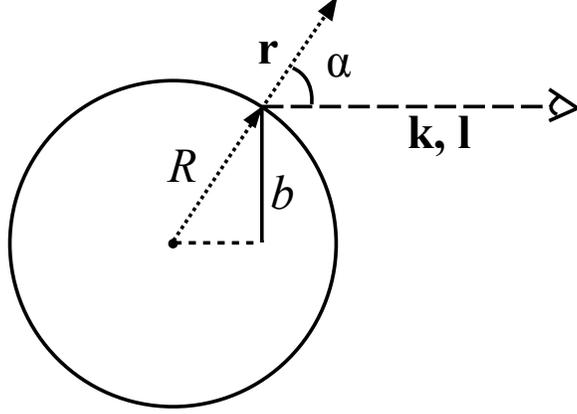


Figure 2.5: Shown in Newtonian gravity, with no light bending (for simplicity). The photon impact parameter b , is the length of the solid line perpendicular to the observer’s line of sight and the equator of the neutron star. Since this is in Newtonian space, we have $\mathbf{k} = \mathbf{l}$, so the photon is emitted in the direction of the observer. In addition, note that R , b , \mathbf{r} , and \mathbf{k} intersect at the photon emission point on the hotspot on the surface of the neutron star.

correct relativistic relationship between α and b as [24]

$$\sin \alpha = \frac{b}{R} \sqrt{1 - \frac{2M}{R}}. \quad (2.2.5)$$

A simpler approximation to the full bending angle integral (equation (2.2.4)) is the Beloborodov bending angle approximation [56],

$$1 - \cos \alpha = (1 - \cos \psi) \left(1 - \frac{2M}{R}\right). \quad (2.2.6)$$

This is only appropriate for small values of $\frac{M}{R}$ and for photon emission close to zenith, since the approximation is purely the Newtonian limit with a relativistic correction factor. We find that while this is a commonly-made approximation in the early literature, the computation of the ray-tracing integral is not too computationally expensive. We thus elect to use the complete formalism to achieve better precision.

2.3 Time of Arrival

As the emitted photons travel from the neutron star's surface to the observer at infinity, an amount of coordinate time (that is, time as experienced by the photon) has passed minus that experienced by the fiducial photon emitted with $b = \alpha = 0$. We compute this amount with the equation [59]

$$T = \int_R^\infty \frac{1}{1 - \frac{2M}{r}} \left\{ \left[\left(1 - \frac{b^2}{r^2} \left(1 - \frac{2M}{r} \right) \right)^{-1/2} - 1 \right] \right\} dr. \quad (2.3.1)$$

We use this equation to reconstruct the photon emission time, since two photons emitted simultaneously but at different phases or impact parameters or local emission angle or in a different initial direction will arrive at the observer at different times [58]. The photons are time-binned relative to when they were received, but this computation allows us to re-bin the photons with respect to when they were emitted. In other words, this computation lines up the photon emission bins and the photon detection bins.

2.4 Doppler Effects

An emitting source moving at a substantial fraction of the speed of light towards the observer will have a bluer observed emission spectrum than the same source in the rest frame. Conversely, if this same source is moving away from the observer at a measurable fraction of the speed of light, the observer will measure the spectrum to be redder than the source at rest. This blueshifting and redshifting of emission spectra is first-order Doppler shifting.

Additionally we account for the transverse Doppler effect. This second-order Doppler shifting arises from motion perpendicular to the line of sight that is present in our system, causing changes in frequency and angles. Doppler effects are required because the neutron stars we model are rotating so rapidly that they can reach relativistic speeds on the surface, $\sim \frac{1}{10}c$. These Doppler

effects, both first order and second order, have large effects on the emission observed at infinity.

The full Doppler boost function, η , is

$$\eta = \frac{\sqrt{1 - v^2}}{1 - v \cos \xi}, \quad (2.4.1)$$

where the speed at the hotspot on the surface of the neutron star v as a fraction of the speed of light is given by [55][59]

$$v = \frac{2\pi\nu_* R \sin \theta}{\sqrt{1 - 2M/R}} \quad (2.4.2)$$

and ξ is the angle between the photon's velocity vector and the photon's initial direction in the observer's frame, given by [55]

$$\cos \xi = -\frac{\sin \alpha \sin i \sin \phi}{\sin \psi}. \quad (2.4.3)$$

Here, ν_* is the neutron star spin frequency in Hz as measured by an observer at infinity, and α , ψ , ϕ , i , and θ are as previously defined. The factor $\sqrt{1 - 2M/R}$ accounts for gravitational time dilation, so that the v calculated here is local, since the ν_* measured is at infinity.

The Doppler boost factor η is highly dependent on phase. It will show redshifting at $\phi \sim \pi/2$ and blueshifting at $\phi \sim 3\pi/2$. It is well-understood that Doppler boosting will cause an aberration in angle [71]. We show here how it affects angles and solid angles in our geometry. The Doppler aberration of the zenith angle α is [71]

$$\cos \alpha_{em} = \eta \cos \alpha_{obs}. \quad (2.4.4)$$

The solid angle of the hotspot, as seen by an observer at a distance D if

the star was not rotating, is given by [24][59]

$$d\Omega_s = \frac{dS_{em}}{D^2} \frac{1}{1 - 2M/R} \cos \alpha \frac{d \cos \alpha}{d \cos \psi}, \quad (2.4.5)$$

where dS_{em} is the surface area for a circular hotspot on a spherical neutron star. This area is [72]

$$dS_{em} = 2\pi R^2 (1 - \cos \rho), \quad (2.4.6)$$

where ρ is the angular radius of the hotspot as previously defined. For a rapidly rotating neutron star, the Doppler aberration modifies equation (2.4.5) so that [73]

$$d\Omega_{obs} = \eta d\Omega_s \quad (2.4.7)$$

is the total solid angle of the hotspot as seen by an observer at a distance D . Note that as η and $d\Omega_s$ change with phase, so does $d\Omega_{obs}$.

Since a neutron star has a large gravitational potential, the emitted photons leaving that potential are somewhat redshifted, which is referred to as “gravitational redshift” to distinguish it from cosmological redshifting of emitting sources at enormous distances. Gravitational redshift is given by

$$1 + z = (1 - 2M/R)^{-1/2}. \quad (2.4.8)$$

Therefore light that is emitted at a frequency ν_{em} is observed at [59]

$$\nu_{obs} = \frac{\eta}{1 + z} \nu_{em}. \quad (2.4.9)$$

The formula for Doppler-boosted intensity is [52][70]

$$I_{obs} = \left(\frac{\nu_{obs}}{\nu_{em}} \right)^3 I_{em}. \quad (2.4.10)$$

2.5 Blackbody Radiation

The observed flux is the intensity over the solid angle of emission, given by [59]

$$dF_{obs} = I_{obs} d\Omega_{obs} . \quad (2.5.1)$$

This produces the total flux as measured by an observer at infinity

$$F_{\nu_{obs}} = \int_{\cos \alpha_{em} > 0} \frac{\eta^4}{(1+z)^3} I_{em} d\Omega_s , \quad (2.5.2)$$

which is integrated over all possible emitted angles. Three powers of η come from the intensity transformation and the fourth comes from the solid angle transformation.

Since we assume completely thermal emission from the hotspot, we compute the flux as a blackbody. The bolometric blackbody flux of the hotspot in the neutron star's rest frame (i.e., non-relativistic) is

$$F_{em} = \sigma T^4 . \quad (2.5.3)$$

The monochromatic flux per phase bin i , in photons $\text{cm}^{-2} \text{s}^{-1} \text{keV}^{-1}$, is given by

$$F_{mono, obs, i} = \frac{1}{E} \cdot \frac{2 \Upsilon d\Omega_{s, i} \eta_i^4}{c^2 h^3 (1+z)^3} \cdot \frac{(E(1+z)/\eta)^3 \times 10^9}{e^{E(1+z)/\eta T} - 1} \quad (2.5.4)$$

where E is the monochromatic energy in keV as observed by an observer at infinity and T is the temperature of the hotspot in keV. The surface area of the hotspot, $d\Omega_{s, i}$, is defined in equation (2.4.5) per phase bin. The graybody factor, Υ , is set to 1 for blackbody, and is explained in section 3.1.2. The redshift factor, $1+z$, is given in equation (2.4.8). The Doppler boost factor, η_i , is defined in equation (2.4.1) per phase bin.

Integrating over all wavelengths, the bolometric flux per phase bin, in pho-

ton count units (photons $\text{cm}^{-2} \text{s}^{-1}$):

$$F_{\text{bolo, obs}, i} = \frac{\Upsilon d\Omega_{s, i} \eta_i^4}{(1+z)^3} \cdot \frac{2 \times 10^9 \pi T^3}{c^2 h^2} \cdot 2.404, \quad (2.5.5)$$

where the last term comes from evaluating a units-modified version of Bradt equation (6.17) [52] using a Riemann zeta function.

2.6 Oblate Approximation

An oblate neutron star is fatter around the middle than over the top, so that the radius at the equator is greater than the radius along the spin axis. Oblateness is an effect of centrifugal forces from rapid rotation, where the difference between light curves from spherical and oblate neutron stars becomes distinguishable at ~ 350 Hz, and is appreciable at $\gtrsim 600$ Hz.

The formalism for dealing with an oblate neutron star model is in Morsink et al. 2007 [59]. For most cases in this thesis we use the spherical model but the light curve simulation code incredibly easily adopts the oblate shape. We use an oblate shape for Cases 6 and 7 in Chapter 6.

One of the key differences between an oblate shape and a spherical shape is that for an oblate neutron star, the normal vector from the surface and the radial vector from the surface are not the same. So we introduce an angle β , the oblate zenith angle, which is defined as

$$\cos \beta = \mathbf{l} \cdot \mathbf{n} \quad (2.6.1)$$

where \mathbf{l} is the direction in which the photon is emitted, and \mathbf{n} is the normal vector to the surface at the hotspot. The full exposition can be found in [59].

The mathematics shown in this chapter provide the structure for the light curve simulation code shown in the next chapter.

Chapter 3

Light Curve Simulation Code

In this chapter we tackle the “forward problem”: we know all the parameters of the neutron star and hotspot, and create the corresponding light curve. This is accomplished using *Spot*, a light curve simulation code. With the input parameters, *Spot* computes the flux in phase bins over one rotation of the neutron star. We then output the flux into a table consisting of the normalized phase, bolometric flux, monochromatic flux at 2 keV, a low energy band flux, and a high energy band flux.

The first part of the chapter discusses hotspot parameters used by the code. The second part of the chapter explains how the code handles non-trivially-sized hotspots. The third part of the chapter discusses the calculation of flux observed in discrete energy bands. The fourth part of the chapter includes how the code was expanded to allow for differently shaped hotspots. The fifth part of the chapter accounts for a non-zero background temperature on the surface of a neutron star during an X-ray burst. The sixth part of the chapter shows the light curves of two antipodal hotspots viewed from different inclination angles. The seventh part of the chapter discusses how the code handles hotspots that cover the rotation axis.

The research in this chapter is a continuation of the *Spot* code written by C. Cadeau and S. Morsink [58][59]. Sections 3.2–3.6 are novel research

completed for this thesis.

For all equations shown, we assume $G = c = 1$.

3.1 Hotspot Parameters

The parameters used to describe a neutron star and the emitting hotspot can be separated into two categories, essential and additional. The essential parameters are required for a neutron star model; the additional parameters are optional, with default values or settings to be used unless otherwise specified. Here we explain all parameters used by the code.

3.1.1 Essential Parameters

The mass, M , is the total mass of the neutron star. When interfacing with the user, the code uses values normalized to the mass of the Sun, with units M_{\odot} . Our nominal value is $1.6 M_{\odot}$.

The radius of the neutron star, R , is measured from the center of the neutron star to the surface at the center of the hotspot, in the rest frame of the neutron star. For spherical neutron stars, this is the same as the radius as measured at the equator, but for oblate neutron stars this is an important distinction. Our nominal value is 12 km.

The spin frequency of the neutron star, ν_* , is the measure of how rapidly the neutron star is rotating about its rotational axis, as observed by an inertial observer at infinity. Since spin frequency can affect the mass and radius values predicted by an equation of state [10], it is very important to account for this. Our nominal value is 581 Hz, which is the observed spin frequency of 4U 1636–536, a known thermonuclear burst oscillation source.

The inclination angle, i , shown in Figure 2.1, is defined as the angle between the rotation axis of the neutron star and the observer’s line of sight. In Figure 3.1 we observe how changing the inclination angle affects the light curve. All

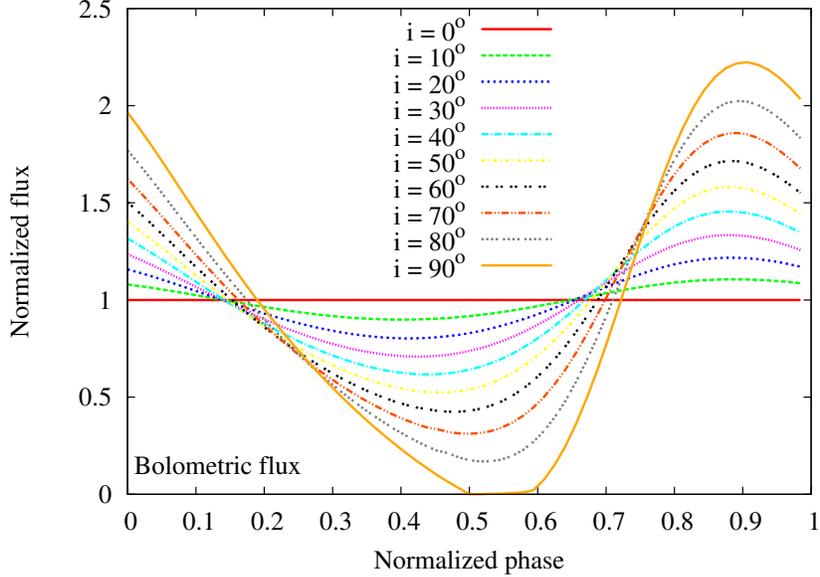


Figure 3.1: Light curves for a spherical neutron star with $M = 1.6 M_{\odot}$, $R = 12 \text{ km}$, varying inclination angle, $\theta = 40^{\circ}$, $\rho = 6^{\circ}$, and $\nu_{\star} = 581 \text{ Hz}$. The key indicates the inclination angle of each light curve. For example, the solid orange curve is the light curve for a neutron star with the listed parameters and an inclination angle of 90° .

parameters are fixed except for the inclination angle, which is allowed to range from 0° to 90° in increments of 10° . Note that for an inclination angle of 0° , the light curve is a straight line at a normalized flux of 1. This is because at an inclination of 0° , an observer is looking straight down the pole of the neutron star. In this scenario, the rotation of the neutron star does not cause the hotspot to appear bigger or smaller; instead, the hotspot is either entirely visible or, if the hotspot is located close to the opposite pole, never visible. Since the same surface area of the hotspot is always visible, we receive the same amount of photons from the hotspot at all points in one rotation period, and thus we do not see any modulation in the light curve.

The emission angle, θ , is measured at the center of the neutron star between the radial vector at the rotation axis and the radial vector going through the center of the hotspot (Figure 2.2). In Figure 3.2 all parameters are fixed except for the emission angle, which is allowed to range from 0° to 90° in increments

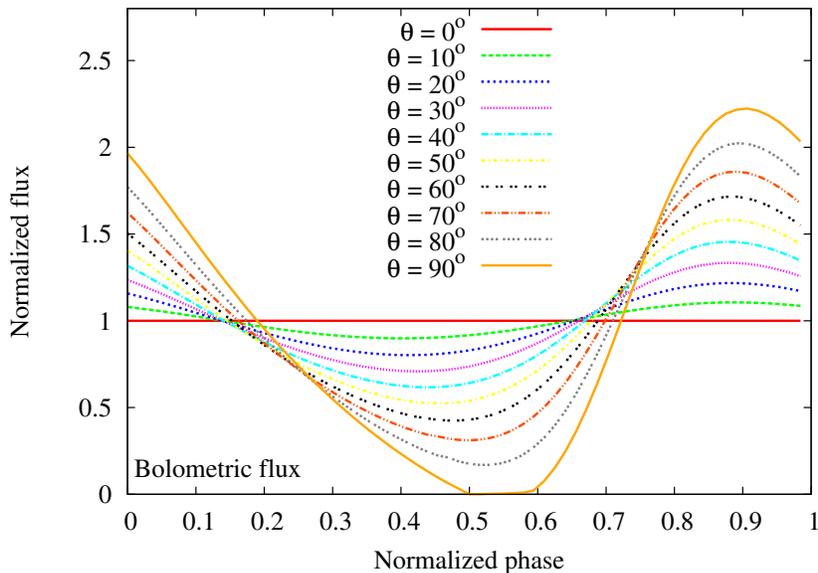


Figure 3.2: Light curve for a spherical neutron star with $M = 1.6 M_{\odot}$, $R = 12 \text{ km}$, $i = 40^{\circ}$, varying emission angle, $\rho = 6^{\circ}$, and $\nu_{\star} = 581 \text{ Hz}$. These are the same parameters as in Figure 3.1, except we fix the inclination angle at 40° and allow emission angle to vary. We notice that this plot looks like a carbon copy of Figure 3.1. See discussion in text.

of 10° . As with inclination, an emission angle of 0° gives a straight line at a normalized flux of 1. This is because at an emission angle of 0° , the hotspot sits symmetrically over the geographic pole. In this scenario, the rotation of the neutron star does not cause the hotspot to move, and so its size is constant over one phase. Since the same surface area of the hotspot is always visible, we receive the same amount of photons from the hotspot at all points in one rotation period, and thus we do not see any modulation in the light curve.

The plots in Figures 3.1 and 3.2 are identical. From this we infer that the inclination angle and emission angle play very similar roles in the appearance of a light curve, and can be considered degenerate parameters. This important point will be explored further in Chapter 5.

The angular radius of the hotspot, ρ , is measured at the center of the neutron star between the radial vector pointing to the center of the hotspot and the radial vector pointing to the outer edge of the hotspot (Figure 2.2).

For light curves shown in this and subsequent chapters we adapt $\rho = 6^\circ$, unless otherwise specified in the figure caption. Section 3.2 discusses how hotspot size affects light curve.

The phase angle, ϕ , is not specified by the user, but is important in understanding the functionality of the code. The user can specify how many phase bins they want the code to use when computing the flux for the light curve. The code then divides this number into 2π (since 2π is one phase, or one rotation of the neutron star) to get the amount of phase (in radians) represented by each bin. Note that for all plots shown in this thesis we normalize the phase period to 1. We use the “front” of each phase bin as the phase value for its corresponding flux, such that normalized phase ranges from 0 to $1 - \textit{phase bin}$. The code computes the flux for each phase bin, and each phase bin is one data point on the light curve. All light curves shown are actually a series of data points that are plotted with lines for better visualization. All light curves begin with the hotspot at $\phi = 0$, and the observer has a constant position at $\phi = 0$.

The *Spot* code has the capability to model two types of neutron stars, oblate and spherical, as described in Chapter 2. The desired shape is set by the user with a simple toggle switch.

We expanded the code’s flux-calculating ability to include computing the flux over different energy bands. The upper and lower limits on these bands can be specified by the user at the command line. See Section 3.3 for details.

3.1.2 Additional Parameters

Since there is a time delay from when the photons are emitted to when they arrive at a detector, we re-bin the photons in accordance with their calculated relative arrival time to re-construct the emitted light curve from the detected light curve as described in Section 2.3 [59].

The user can specify an output directory and file name. This is the file where the table of fluxes are written for the given neutron star parameters. We

use a directory naming convention that makes the parameter values part of the folder and file names, so that it is easy to classify and find the output tables.

We can set the distance from the observer to the neutron star. All models in this thesis use a distance of 6 kpc. Since changing the distance will only affect the solid angle of the spot $d\Omega_{obs}$, and we normalize the flux to 1, changing the distance does not affect the shape of the light curve.

The user can toggle the beaming model, either isotropic or graybody. The isotropic model assumes that photons emitted at all zenith angles α (equation (2.2.1)) travel through the same amount of atmosphere and thus have the same opacity. The graybody model is an inefficiently emitting limb-darkening atmosphere (independent of wavelength) such that photons emitted close to the normal to the surface travel through a less opaque total atmosphere, and photons emitted close to tangent to the surface travel through a more opaque total atmosphere [52]. This makes photon emission perpendicular to the surface brighter than photon emission tangent to the surface. Our code uses a Hopf function to approximate the graybody [74]. We typically use isotropic flux, but two cases in Chapter 6 have a graybody flux.

The user can designate a data file to compare with the model parameters they specify. The code computes the χ^2 fit to compare how well the specified light curve model fits the data using equation (4.0.1). This function is a crucial aspect of Chapters 4, 5, and 6.

A phase shift is used when comparing a model with data. This value shifts the light curve in the x -direction. A light curve can “start” at any point in the phase, so this parameter lines up our model with the data so that they are in phase for direct comparison. Since we use a phase normalized to 1, the phase shift is also normalized. We use a phase shift of 0 for all simulated data in this thesis. Note that this parameter is circular in the range 0–1, so 0.99 is effectively very close to 0.00.

The user is able to set a toggle flag to normalize the flux. The code computes

the flux in photon count units (photons $\text{cm}^{-2} \text{s}^{-1}$), but we typically normalize the flux so that the mean value of the light curve is 1.

For a large hotspot, we can compute the flux for smaller surface area chunks and sum the values per phase bin. This is practical for hotspots large enough to span a range of emission and phase angles where evaluating the flux only at the center of the hotspot is not a good approximation. See Section 3.2 for details.

Another optional parameter is the kT temperature of the hotspot, in units of keV, as measured in the neutron star’s rest frame. For all models presented in this thesis we assume a hotspot temperature of 2 keV, consistent with Kuulkers et al. 2003 [75], Bhattacharyya et al. 2005 [45], and Lo et al. 2013 [39]. We treat the ambient surface temperature of the neutron star to be 0 keV as in Braje et al. 2000 [76], since it is likely significantly cooler than the hotspot. We discuss the background temperature in further detail in section 3.5. The k in kT is dropped in later references to temperature, but it is implied.

We can input a temperature mesh to achieve different hotspot shapes. The effects of this are addressed in Section 3.4. The standard hotspot shape, used for all cases shown in this thesis, is a full circle.

Calculating the flux from two antipodal hotspots on the surface of the neutron star can be toggled with a flag. See Section 3.6 for details. We are also able to compute the flux from only the second antipodal hotspot by setting a different flag. This is useful when debugging and for demonstrating how the top hotspot and second hotspot add up to make the components of the two-hotspot light curve.

3.2 Non-Trivially-Sized Hotspots

The parameter ρ , the angular radius of the hotspot, allows us to change the size of the hotspot (Figure 2.2). However it was unclear if the code is as precise

for big hotspots as it is for infinitesimal hotspots. To accurately accommodate large hotspots we split them up into a mesh or grid overlaid on the hotspot and compute the flux for each cell, summing the values within each phase bin. All meshes consist of $n \times n$ cells covering the circular hotspot. We show an example of a 4×4 mesh with approximate cell shapes in Figure 3.3. A 1×1 mesh is geometrically equivalent to treating the whole hotspot as one uniformly emitting region.

The code steps through the mesh beginning with the top leftmost cell. The code finds the local emission angle,

$$\theta_k = \theta - \rho + k \cdot d\theta + 0.5 \cdot d\theta, \quad (3.2.1)$$

where k is the row index ranging from 0 to $n-1$ and $d\theta = 2\rho/n$. Using spherical trigonometry, the phase angle at the edge of the current row is

$$\phi_{\text{edge}} = \cos^{-1} \left(\frac{\cos \rho - \cos \theta \cos \theta_k}{\sin \theta \sin \theta_k} \right), \quad (3.2.2)$$

assuming that the center of the hotspot is at $\phi = 0$. We then compute the local phase angle for that specific cell,

$$\phi_j = -\phi_{\text{edge}} + j \cdot d\phi + 0.5 \cdot d\phi, \quad (3.2.3)$$

where j is the column index ranging from 0 to $n-1$ and $d\phi = 2\phi_{\text{edge}}/n$. The surface area for cell (k, j) of angular area $d\theta d\phi$ is approximated as

$$dS_{em, (k, j)} = R^2 \sin |\theta_k| d\theta d\phi \quad (3.2.4)$$

so that the total hotspot surface area is

$$dS_{em} = \sum_{k=0}^{n-1} \sum_{j=0}^{n-1} dS_{em, (k, j)} \quad (3.2.5)$$

instead of equation (2.4.6). This local θ_k , ϕ_j , and $dS_{em,(k,j)}$ are then used in the equations in Chapter 2 to compute the flux of photons emitted from each cell.

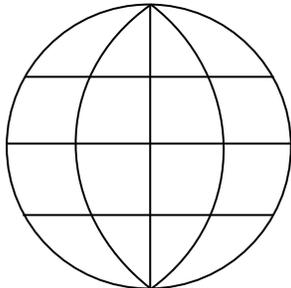


Figure 3.3: Sample 4×4 mesh over a hotspot. The hotspot is the outlined circle, and this mesh splits up the hotspot into 16 total cells. Increments in ϕ move in the x -direction, and increments in θ move in the y -direction. Within *Spot* each cell is assigned the local θ and ϕ that are appropriately offset from the hotspot's overall θ and ϕ , and flux is calculated for each cell per phase bin. Then, for each phase bin, we add up the flux from each cell to get the total flux from the neutron star within each phase bin.

We model a smaller hotspot, with an angular radius of 6° , in Figure 3.4. We notice that dividing the hotspot into a mesh and calculating fluxes for each mesh bit individually does not affect the overall flux. Thus we do not find a significant difference due to mesh finenesses. We then model a light curve from a larger hotspot, with an angular radius of 20° , in Figure 3.5. While the 1×1 mesh light curve is distinct, there is no discernible difference between the light curves using the finer meshes.

We find that the light curves of larger hotspots are less modulated (have a smaller pulse amplitude) than the light curves of smaller hotspots. This is because the larger hotspot is visible for a wider range of phase angles, and the larger hotspot is more strongly lensed on the opposite side of the neutron star. This gives a less noticeable difference between the bright and dim parts of the phase, and the light curve is dim for a smaller fraction of the phase. Thus we observe a smaller pulse amplitude. The differences between the light curves of different meshes are much more pronounced for large hotspots. In Figures 3.4

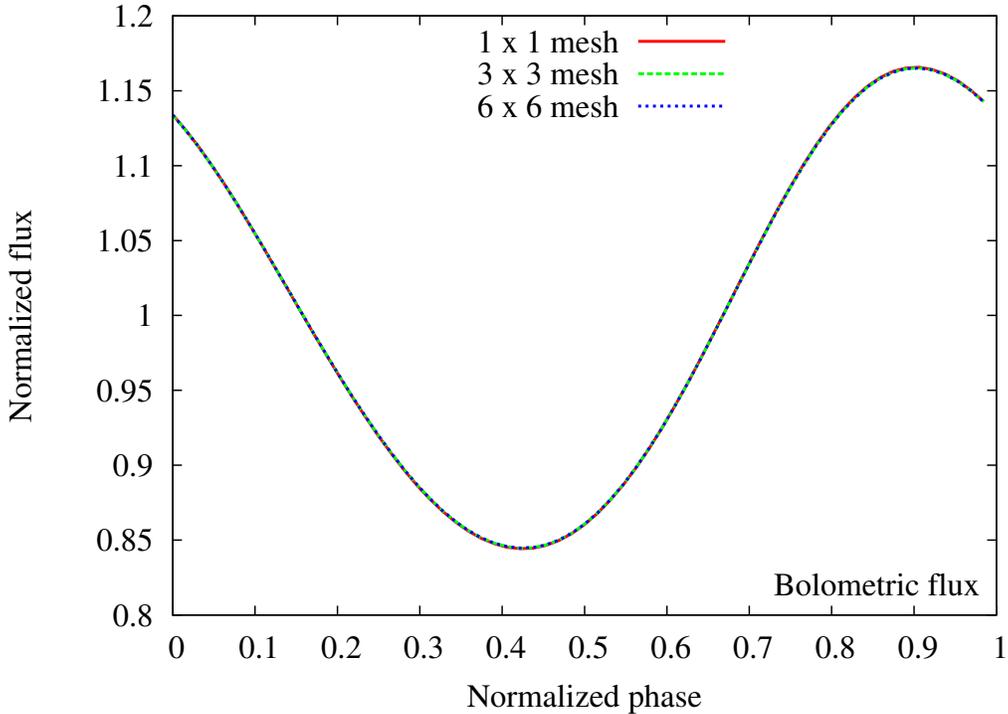


Figure 3.4: Light curves for a spherical neutron star with $M = 1.5 M_{\odot}$, $R = 12 \text{ km}$, $i = 20^{\circ}$, $\theta = 40^{\circ}$, $\rho = 6^{\circ}$, $\nu_{\star} = 581 \text{ Hz}$, and varying mesh fineness. We see that there is no noticeable difference between the light curves.

and 3.5, we notice that when the size of each cell is less than approximately 6° in radius, there is no noticeable difference between the light curves.

Additionally, a finer hotspot mesh is computationally expensive compared to the time of one full light curve computation for a 1×1 mesh. Therefore in Chapter 6 we use a hotspot radius of 6° with a 1×1 mesh.

3.3 Flux in Energy Bands

To improve the applicability of our code to observations from telescopes, we expand its capabilities to include calculating flux for specific energy bands. This will allow the code to fit data from telescopes and thus be more realistic, since telescopes do not detect bolometric or monochromatic flux. The equation

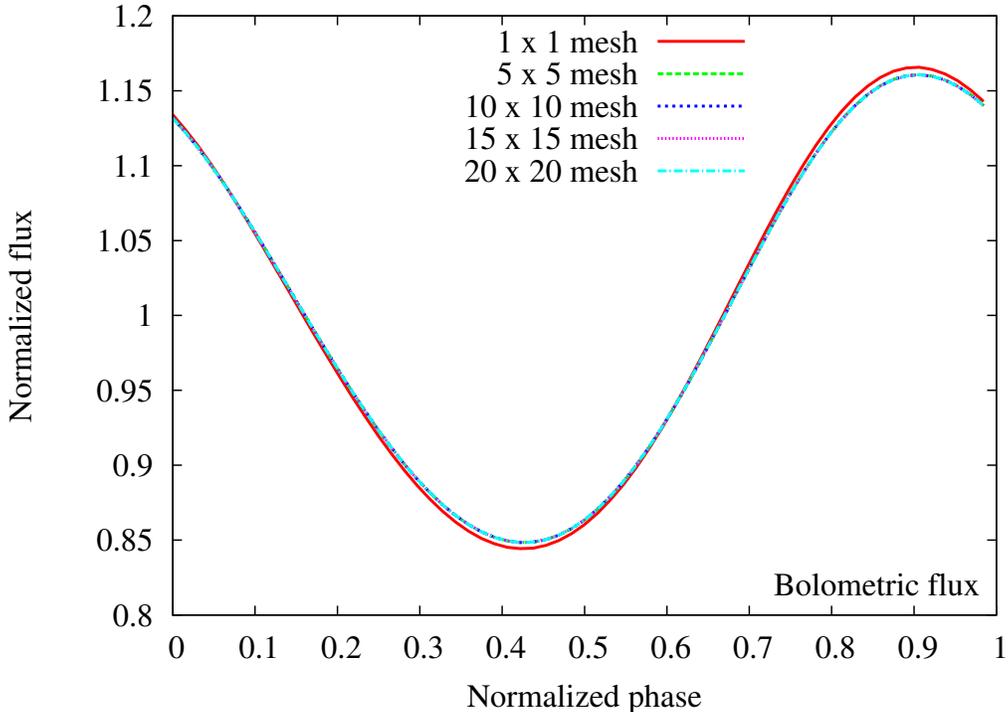


Figure 3.5: Same parameters as Figure 3.4 but with $\rho = 20^\circ$, and accordingly different meshes. We notice that the 1×1 mesh light curve is distinct from the rest.

for calculating flux in an energy band, in photon count units, is [52]

$$F_{obs,i} = \frac{2T^3}{c^2h^2} \frac{\Upsilon \eta_i^4 d\Omega_i}{(1+z)^3} \int_{E_1/T}^{E_2/T} \frac{x^2}{e^x - 1} dx. \quad (3.3.1)$$

Our integration variable is $x = E/T$, so that $dx = dE/T$. The integration bounds are from $E_1 = E_{\text{lower}} \frac{1+z}{\eta}$ to $E_2 = E_{\text{upper}} \frac{1+z}{\eta}$, where E_{upper} is the upper bound of the energy band in keV as measured by an observer at infinity, and E_{lower} is the lower bound of the energy band in keV, also as measured by an observer at infinity. T is the temperature of the hotspot in keV and Υ is the graybody factor. The surface area of the hotspot, $d\Omega_i$, is defined in equation (2.4.5) per phase bin. The redshift factor, $1+z$, is given in equation (2.4.8). The Doppler boost factor, η_i , is defined in equation (2.4.1) per phase bin. Since the Doppler boost factor changes with the phase angle ϕ , the integration limits change with

phase as well. F_i is the total flux as measured in each phase bin.

To calculate the flux in units of $\text{erg cm}^{-2} \text{s}^{-1}$, we have

$$F = \frac{2T^4}{c^2 h^3} \frac{\Upsilon \eta^4 d\Omega}{(1+z)^3} \int_{E_1/T}^{E_2/T} \frac{x^3}{e^x - 1} dx, \quad (3.3.2)$$

where all variables are the same as in equation (3.3.1).

As an example, we calculate what the flux for a neutron star model would look like in energy bands detectable with past, current, and proposed X-ray telescopes in Figure 3.6, assuming 100% flat detector efficiencies. These light curves have the same parameters but are calculated for different energy bands.

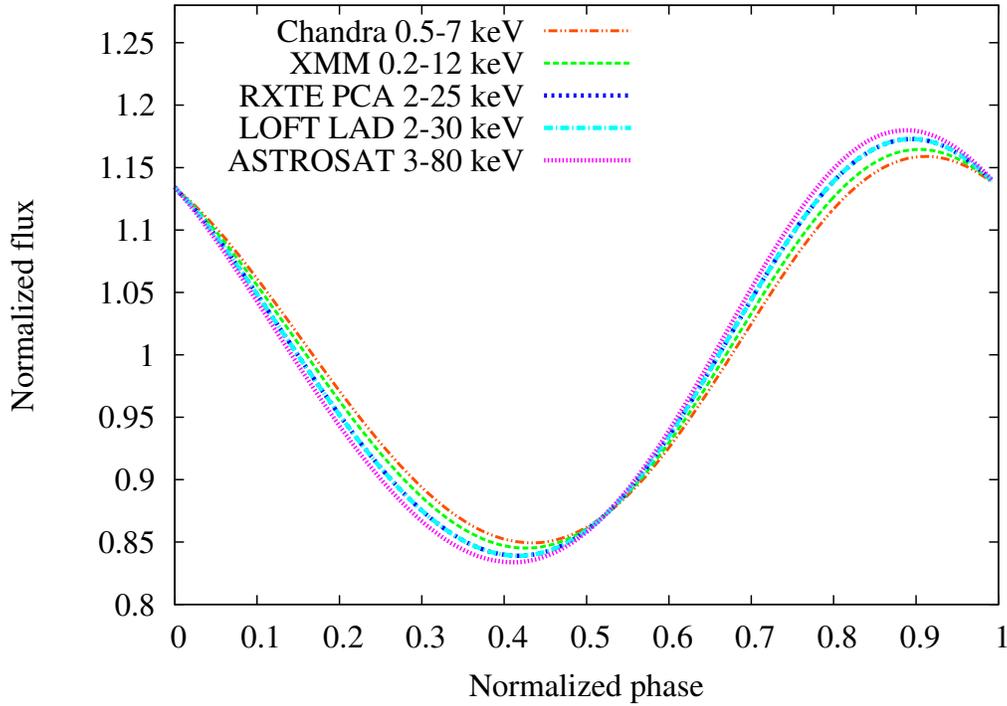


Figure 3.6: Light curves as they would appear in various telescope instrument energy bands for a spherical neutron star with $M = 1.6 M_{\odot}$, $R = 12 \text{ km}$, $i = 10^{\circ}$, $\theta = 60^{\circ}$, $\rho = 6^{\circ}$, and $\nu_{*} = 581 \text{ Hz}$. Note that the RXTE and LOFT curves effectively overlap. The energy band values are from [61][77] [78][79][80].

Since we see that the energy ranges 2–3 keV and 5–6 keV are covered by all

of the different telescopes, we choose to use these as the two standard energy bands computed by our code for all models in Chapter 6. These are shown in Figure 3.7. While we cannot detect a true bolometric flux from a telescope, we include it as a baseline for the code.

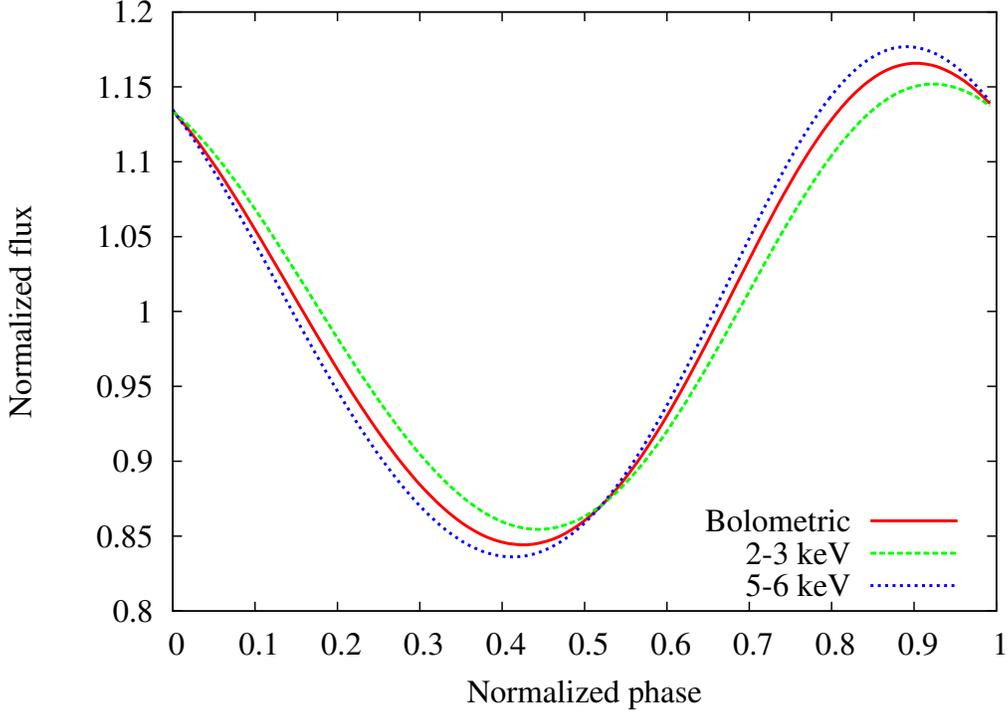


Figure 3.7: Light curves for a spherical neutron star with $M = 1.6 M_{\odot}$, $R = 12 \text{ km}$, $i = 10^{\circ}$, $\theta = 60^{\circ}$, $\rho = 6^{\circ}$, and $\nu_{*} = 581 \text{ Hz}$. Each curve here represents the light curve flux in the labeled energy range. The soft lag is apparent (see discussion).

We notice that the higher energy band light curve peaks earlier in the phase than the low energy band light curve. Furthermore, the low energy curve has a smaller pulse amplitude than the high energy curve. These two effects are due to the Doppler boosting of the photons. Since blueshifted Doppler boosting will increase the energy of the photons, the photons experience a maximal blueshift at the peak Doppler boost factor. This occurs at $\phi_{norm} = 0.75$, when the hotspot is tangential to our line of sight and moving towards us. So more energetic photons will peak at a normalized phase close to 0.75 and less energetic photons

will peak at a later phase. However we also have to account for the solid angle of the hotspot as apparent to the observer at infinity. This is much bigger at phases past 0.75 than phases before, so the energetic photons will peak at a phase close to but past $\phi_{norm} \sim 0.75$ and less energetic photons will peak at a phase after that. Thus, the peak flux for a detected energy of 6 keV occurs a little earlier in the phase than for a detected energy of 3 keV. This is part of what gives us the “soft lag” phenomena, in which the soft (low energy) X-ray photons peak later in the phase than the hard (high energy) X-ray photons. Furthermore, the larger the apparent solid angle of the hotspot, the higher the flux is. The combination of these two factors is what gives us the smaller amplitude for soft photons and larger amplitude for hard photons.

Note that the energy bands are for the energy of the photons as detected by an observer at infinity, but the Planck function in Figure 3.8 is in the neutron star’s rest frame from a uniformly emitting 2 keV hotspot. This blackbody curve is Doppler boosted and gravitationally redshifted before arriving at the observer. The redshift will cause the peak of the blackbody curve to shift towards the lower energy end of the spectrum.

To demonstrate the versatility of the code, we test much higher energy bands. In Figure 3.10, we use the ranges 10–50 keV and 50–100 keV. The soft lag is very apparent. In addition, the high energy band curve is very asymmetric. This is due to potent Doppler boosting. We also notice that this light curve peaks in flux quite close to $\phi_{norm} = 0.75$, when the peak Doppler boosting occurs. We hope these extended energy band capabilities prove that our code will be useful when ASTROSAT launches next year.

3.4 Hotspot Shapes

We can change the shape of the hotspot, the region on the surface of the neutron star that is emitting photons. Previous iterations of this code have assumed

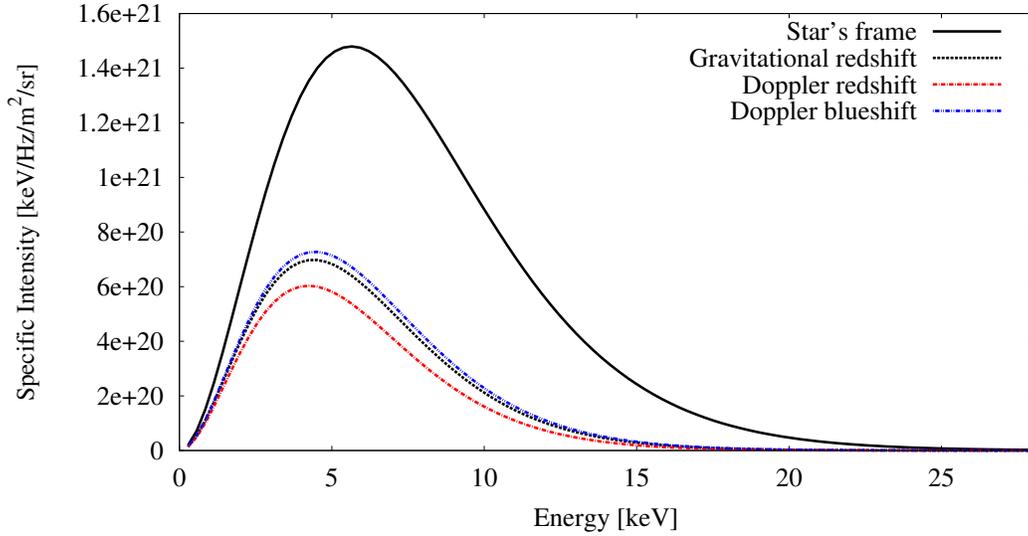


Figure 3.8: Planck function for a $kT = 2$ keV hotspot for the parameters of Chapter 6 Case 1; energy in keV plotted against the specific emitted intensity in units $\text{keV Hz}^{-1} \text{m}^{-2} \text{sr}^{-1}$. The black solid curve is the blackbody spectrum of the hotspot in the neutron star’s rest frame. The black dotted curve includes gravitational redshift, and so would be measured by an observer at infinity if the neutron star was minimally rotating. The red dashed curve includes the maximal Doppler redshift of the rotating neutron star, as measured by an observer at infinity. This occurs at $\phi_{norm} = 0.25$. The blue dashed curve includes the maximal Doppler blueshift of the rotating neutron star, as measured by an observer at infinity. This occurs at $\phi_{norm} = 0.75$. The black solid curve peaks at 5.643 keV, the black dotted curve peaks at 4.393 keV, the red dashed curve peaks at 4.236 keV, and the blue dashed curve peaks at 4.439 keV.

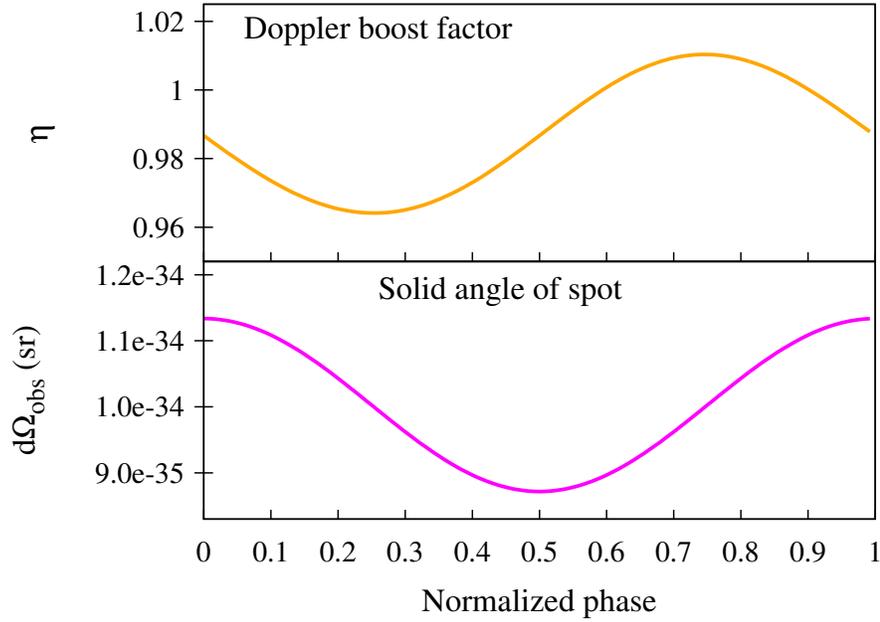


Figure 3.9: A plot of the Doppler boost factor against normalized phase, in orange, and the solid angle of hotspot against normalized phase, in magenta. The Doppler boost factor, η , is calculated using equation (2.4.1), and is unitless. An $\eta > 1$ means blueshifted, $\eta < 1$ means redshifted, and $\eta = 1$ means not shifted. The solid angle of the hotspot as seen by an observer at infinity, $d\Omega_{\text{obs}}$, is calculated using equation (2.4.7). These values have units of steradians. The η and $d\Omega_{\text{obs}}$ values were computed for a neutron star with the same parameters as used in Figures 3.7 and 3.10. The Doppler boost factor curve is centered on ~ 0.986 due to the $\sqrt{1-v^2}$ factor in equation (2.4.1).

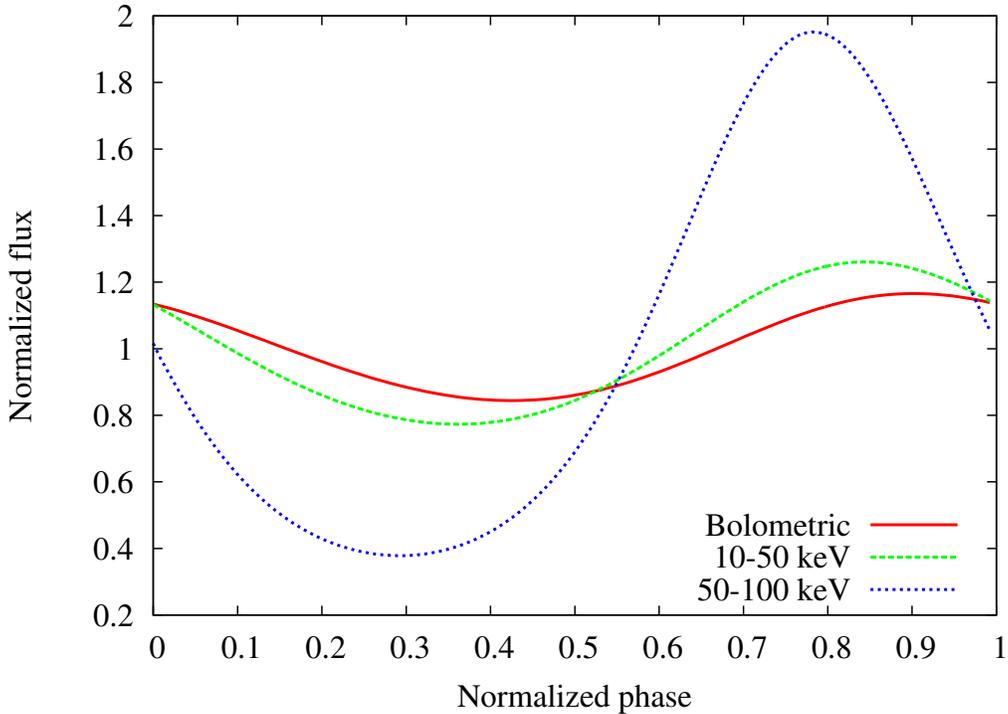


Figure 3.10: Light curves for a spherical neutron star with the same parameters as Figure 3.7. As in Figure 3.7, each curve represents the light curve flux in the labeled energy range. The soft lag is very apparent here (see discussion).

that the hotspot is a solid circle on the surface of the neutron star. While this is the simplest theoretical model (for a non-trivially-sized hotspot), there is no reason to believe that the X-ray burst fuel will burn in an even circular pattern (e.g., see MHD simulations by Kulkarni and Romanova 2005 [51]).

Instead of defining a constant temperature for the whole hotspot, we split the hotspot up into physical bins (as discussed in Section 3.2) and define a temperature for each specific bin. It is important that the number of hotspot mesh bins specified in the command line matches the size of the temperature mesh read in. In this temperature mesh, the mesh bins to be treated as background are assigned a temperature of 0 keV, and the mesh bins to be treated as an emitting region are given a non-zero temperature (e.g., 2 keV). This mesh physically maps to the hotspot, and defines the shape of the region that is emitting photons. A sample temperature mesh is presented in the Appendix,

Table A.1.

We choose to illustrate light curves due to hotspot shapes using two extreme shapes, the superposition of those two shapes, and the full circular hotspot (Figure 3.11). Differences in light curves are much more pronounced for large hotspots with fine meshes. For small hotspots or larger hotspots with very coarse meshes, the differences are minute to negligible.

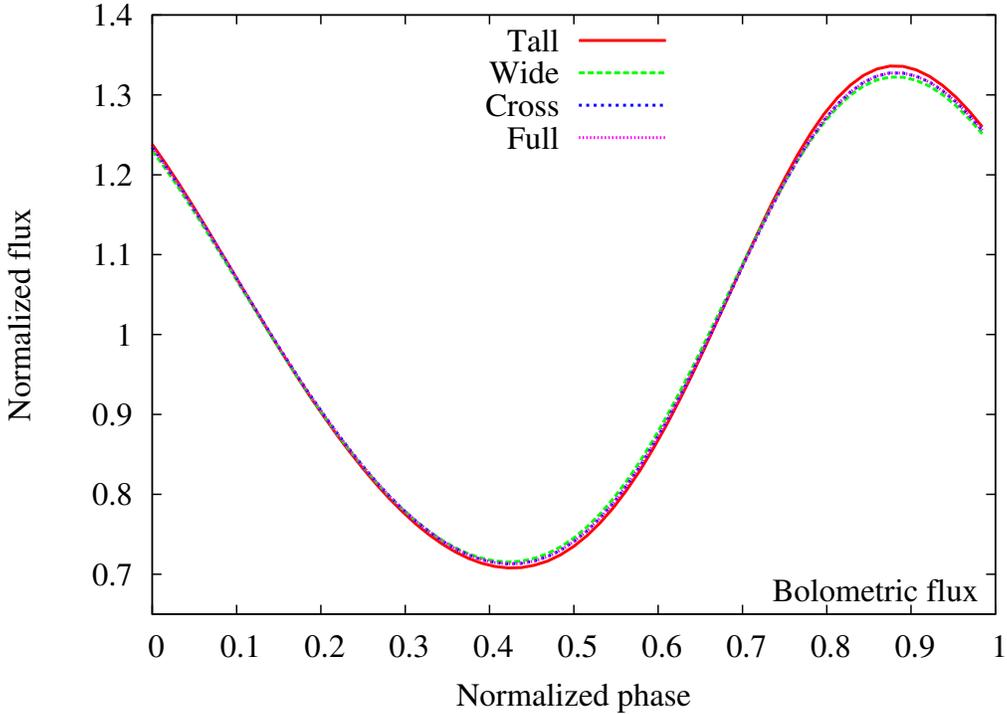


Figure 3.11: Light curves for a spherical neutron star with $M = 1.5 M_{\odot}$, $R = 12 \text{ km}$, $i = 20^{\circ}$, $\theta = 40^{\circ}$, $\rho = 15^{\circ}$, 15×15 hotspot mesh, $\nu_{*} = 581 \text{ Hz}$, and varying hotspot shape. Each curve represents the flux of one period for a neutron star with a hotspot emitting in photons in the specified pattern on the surface of the neutron star. For a visual representation of what each shape means, see Figure 3.12. We notice that “tall” has a slightly bigger pulse amplitude than “full”, “wide” has a slightly smaller pulse amplitude than “full”, and “cross” is not distinguishable from “full”.

The circle referred to in the following descriptions is the outline of a circle with the specified angular radius centered on the specified emission angle, or what constitutes a typical “full” hotspot (see Figure 3.12). “Tall” can be

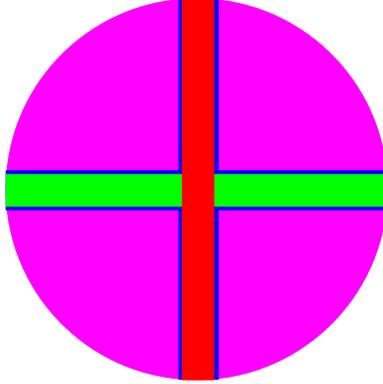


Figure 3.12: A visual representation of the four hotspot shapes we test. “Tall” is the vertical stripe in red, “wide” is the horizontal stripe in green, “cross” is the superposition of vertical and horizontal stripes outlined in blue, and “full” is the entire circle in magenta (including the area covered by the previous shapes). See Figure 3.11 for the light curves for each of these hotspot shapes, given the same neutron star parameters.

thought of as a line of very small ($\rho = 1^\circ$) hotspots spanning a range of emission angles in a thin stripe vertically down the middle of the circle. We know from Figure 3.2 that larger emission angles cause a greater pulse amplitude, because the apparent size of the hotspot varies more during one rotation period. We do indeed see that “tall” has a slightly greater pulse amplitude than “full” (Figure 3.11). “Wide” can be thought of as a line of very small ($\rho = 1^\circ$) hotspots spanning a range of phase angles in a thin stripe horizontally across the middle of the circle. The more phase angles the hotspot encompasses, the less pulse modulation. This is due to the apparent size of the hotspot not changing as much during one rotation period. So as expected, “wide” has a slightly smaller pulse amplitude than “full” (Figure 3.11). “Cross” is a superposition of “tall” and “wide”, such that the emitting region is both vertically down and horizontally across the center of the circle. We theorize that “cross” is indistinguishable from “full” because the increase in pulse amplitude of “tall” approximately cancels the decrease in pulse amplitude of “wide” (Figure 3.11). “Full” is when the entire area of the circle is the emitting region.

From Figure 3.11 we can deduce that emission angle has a bigger effect on

the light curve than the phase angle. We use the full hotspot for all subsequent light curves in this thesis.

3.5 Non-Zero Background Temperature

Accounting for a non-zero background temperature on rest of the surface of a neutron star adds complications to and is very computationally expensive for the *Spot* code. We ran some modified versions of *Spot* accounting for a background surface temperature, and found that it shifted the mean value of the non-normalized light curve in the positive y -direction, but did not change the non-normalized pulse shape. This means that normalizing the shifted light curve will change the overall pulse shape by a minor amount (for example, $y = \sin(x)$ and $y = \sin(x)+1$ do not normalize to the same curve). Therefore, the pulse amplitudes computed by *Spot* are lower limits of what would be measured if a background temperature is relevant. It is possible that thermonuclear burst oscillations in the rise of a Type I X-ray burst would have a negligible background temperature (as assumed in Braje et al. 2000 [76] for persistent pulsations).

3.6 Two Antipodal Hotspots

Spot has the ability to calculate the flux from two antipodal hotspots, for example one on the north magnetic pole and one on the south magnetic pole (assuming a simplistic and centered magnetic dipole). This is not strictly necessary to realistically model light curves from Type I X-ray bursts, but it will be important for modelling the light curves from accreting millisecond X-ray pulsars with their larger magnetic fields.

The top hotspot, or first hotspot, is located in the northern hemisphere ($0^\circ \leq \theta \leq 90^\circ$) of the neutron star, and is the same as the single hotspot that

has been shown previously. The bottom hotspot, or second hotspot, is located in the southern hemisphere ($90^\circ \leq \theta \leq 180^\circ$) of the neutron star and at the opposite phase of the first hotspot. So if the first hotspot is face-on to the observer's line of sight at $\phi = 0$, the second hotspot would be face-on to the observer's line of sight at $\phi = \pi$.

For calculating the flux of the second hotspot, we keep the neutron star geometry the same and changed the inclination angle and the phase angle such that $i_2 = 180 - i$ and $\phi_2 = \phi + 180$. We calculated the flux from each hotspot individually, and then added up the flux from both hotspots per phase bin to create the visible light curve.

We also want to explore the effect of inclination on the light curves of two antipodal hotspots. We elect to display the light curves in photon count units instead of a normalized flux, since the normalization only takes into account flux from that particular hotspot and not from both light curves. Thus normalized scales would be misleading. We predict that, since we have shown previously that inclination and emission angles are degenerate with each other, keeping the inclination angle constant while changing the emission angles would produce similar results.

3.6.1 Low Inclination

In this subsection we look at the light curves of a neutron star with two antipodal hotspots, one at 60° and another at 120° , viewed at an inclination angle of 10° from the rotation axis (Figure 3.13).

The dashed line indicates the top hotspot. The second line, dotted, is for the second antipodal hotspot, which is 180° across the neutron star from the top hotspot, as if it were on the opposite end of the magnetic dipole. The solid line is the two light curves added together — this shows what we would actually detect with a telescope if we were looking at this system. We note that for this low-inclination system, the amplitudes are rather small compared

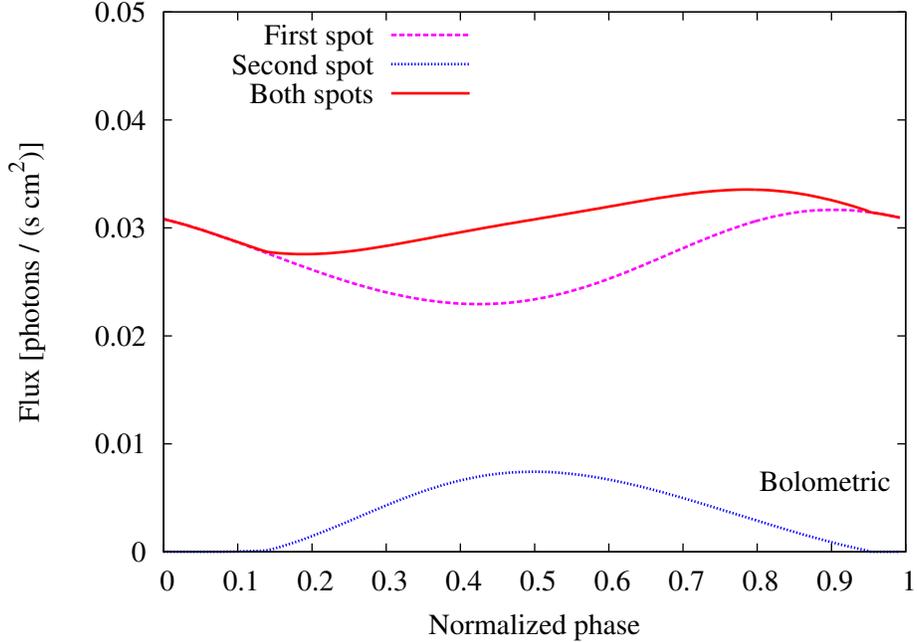


Figure 3.13: Low inclination: Light curves for two antipodal hotspots on a spherical neutron star with $M = 1.6 M_{\odot}$, $R = 12 \text{ km}$, $i = 10^{\circ}$, $\theta = 60^{\circ}$ and 120° , $\rho = 6^{\circ}$ each, and $\nu_{*} = 581 \text{ Hz}$.

to the next two figures.

3.6.2 Moderate Inclination

In this subsection we look at the light curves of a neutron star with two antipodal hotspots with the same geometry as the previous subsection, $\theta_1 = 60^{\circ}$ and $\theta_2 = 120^{\circ}$, but viewed at an inclination angle of 30° (Figure 3.14).

As with Figure 3.13, the dashed line is the first hotspot, the dotted line is the second hotspot, and the solid line is the sum of the two. The second hotspot affects the total curve much more strongly here than in Figure 3.13, creating a distinct bump in the middle. It should be noted that this neutron star has the same parameters as the previous one. The only difference is the viewer's inclination angle to the system. We can see that the total light curve is distinctly non-sinusoidal.

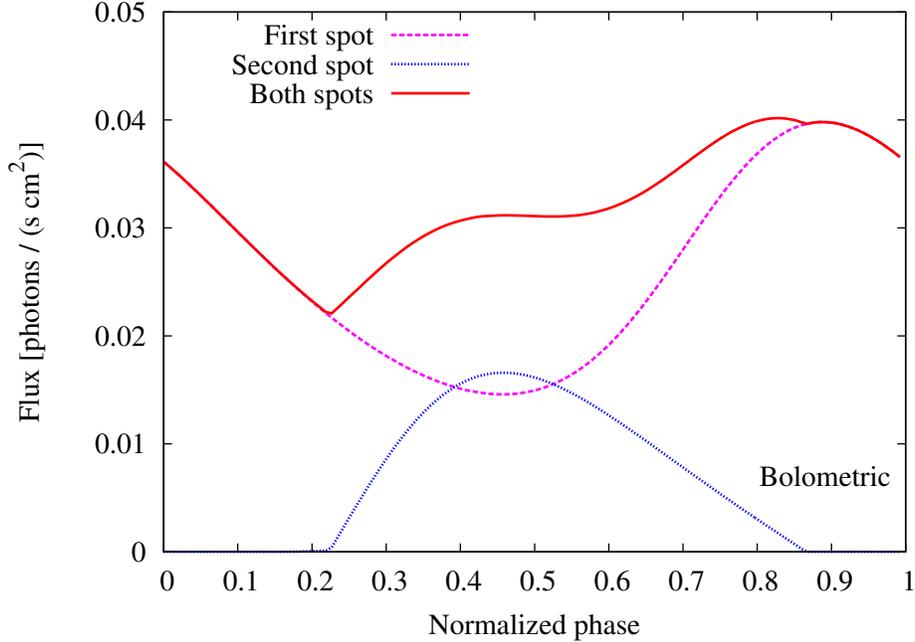


Figure 3.14: Moderate inclination: Light curves for two antipodal hotspots on a spherical neutron star with $M = 1.6 M_{\odot}$, $R = 12 \text{ km}$, $i = 30^{\circ}$, $\theta = 60^{\circ}$ and 120° , $\rho = 6^{\circ}$ each, and $\nu_{*} = 581 \text{ Hz}$.

3.6.3 High Inclination

In this subsection we look at the light curves of a neutron star with two antipodal hotspots with the same geometry as the previous subsections, $\theta_1 = 60^{\circ}$ and $\theta_2 = 120^{\circ}$, but viewed at a large inclination angle of 60° (Figure 3.15).

Here, we very clearly see a middle peak from the second hotspot. In systems with two visible hotspots, we may be able to constrain the inclination angle based on how much each hotspot contributes to the total light curve.

3.7 Hotspot Covering the Rotation Axis

To expand the functionality of the code, we allow for a hotspot covering the rotation axis. We do not think that systems with this geometry will be prevalent in the data, but we want to permit them nonetheless, if only for the sake of completeness. Given that this geometry provides minimal changes to the light

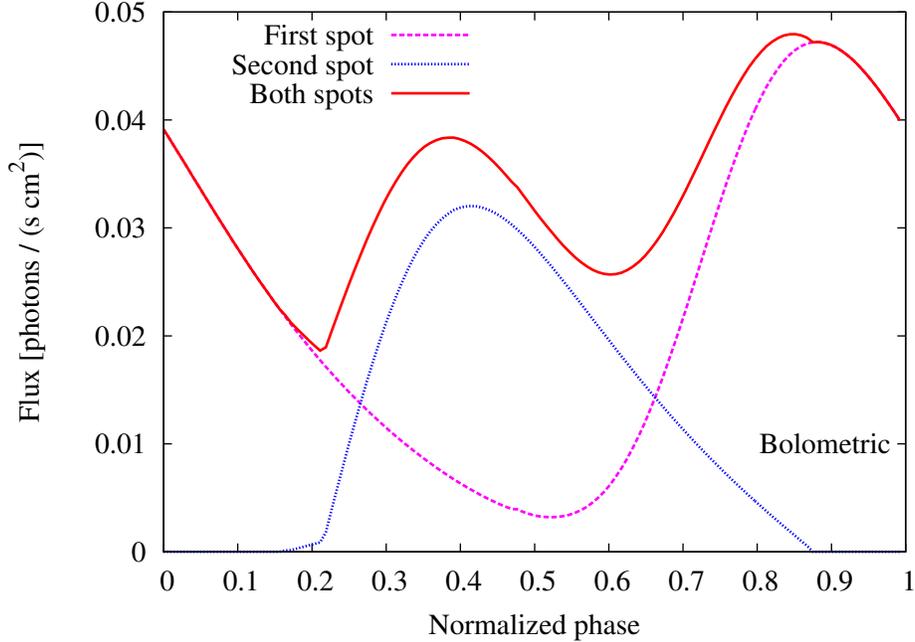


Figure 3.15: High inclination: Light curves for two antipodal hotspots on a spherical neutron star with $M = 1.6 M_{\odot}$, $R = 12 \text{ km}$, $i = 60^{\circ}$, $\theta = 60^{\circ}$ and 120° , $\rho = 6^{\circ}$ each, and $\nu_{*} = 581 \text{ Hz}$.

curve shape, we will not display light curves for the following subsections.

3.7.1 Symmetric Over the Rotation Axis

In this configuration, the hotspot is symmetric over the rotation axis such that $\theta = 0$. As we can see in the first curve of Figure 3.2, this gives us a constant flux that normalizes to 1. Since inclination angle is constrained to the range 0° to 90° , a hotspot centered at $\theta = 0$ will always be at least partially visible to the observer, and it will have the same apparent surface area $d\Omega$ during all times in one complete phase. We compute the flux for $\theta = 0 + \epsilon$, since the equations are undefined at $\theta = 0$, and step through the phase from 0 to 2π in increments of $2\pi / \text{number of phase bins}$.

3.7.2 Asymmetric Over the Rotation Axis

In this configuration, the hotspot is asymmetric over the rotation axis (i.e., $\theta - \rho \leq 0$). This has two sub-cases: computing the flux for the hotspot as a whole, and splitting up the hotspot into multiple bins using a mesh as in Section 3.2. The first case does not require any special treatment since $\theta > 0$, and we compute the flux normally as described in Chapter 2 and Section 3.3.

For the second case, the mesh-and-spherical-trigonometry approach does not hold, so we modify our geometric approach to compute ϕ_j . Instead of slicing up the area of the hotspot into a mesh, we step our way around the circumference of the hotspot. Each circumference segment is assigned the local emission angle θ_k in the same way as before, but with $\phi_j = \phi_{\text{edge}}$ (using equation (3.2.2)). Then, if θ_k is over the pole, we use the absolute value of θ_k but with $\phi_j + \pi$ to evaluate on the “other side” of the pole of the neutron star.

This approximation is acceptable for small hotspots, but is not adequate for large hotspots.

Spot provides the ability to make neutron star models, which we will fit to simulated data using a genetic algorithm as discussed in the next chapter.

Chapter 4

Fitting X-ray Light Curves with a Genetic Algorithm

In this chapter we tackle the backward problem: we take an observed light curve and attempt to determine the parameters of the neutron star and hotspot that created it. This is accomplished by modelling a light curve with a set of parameters as in Chapter 3, and calculating the fit of this curve with the observed data. After testing a multitude of solutions using a genetic algorithm, we select the parameter set which best reproduces the data. All simulated data in this thesis was created by a forward run of the *Spot* code.

We use χ^2 statistics to compare how closely two light curves match. This is computed as

$$\chi^2 = \sum_{i=1}^N \left(\frac{F_{i, \text{data}} - F_{i, \text{fit}}}{\sigma_i} \right)^2 \quad (4.0.1)$$

where σ_i is the error in the i^{th} phase bin (simulated as Poisson noise for an appropriate telescope), $F_{i, \text{data}}$ is the flux from the data file in the i^{th} bin, and $F_{i, \text{fit}}$ is the flux from the fit model for the i^{th} bin. This equation sums over all of the energy bands for the light curve. The lower the χ^2 the better the fit, so the best fit will have the lowest χ^2 value.

With the simulated data used in this chapter, we know the correct set of

parameters that describe the neutron star and hotspot, and are attempting to re-create that with the fits. We will need to perfect the algorithm and do many blind tests that yield accurate results before fitting real data, for which the true set of parameters is not known.

The first section of this chapter explains the function of a genetic algorithm and why it is applicable to our problem. The second section discusses the details of the specific genetic algorithm we are using. The third section discusses two other light curve fitting algorithms seen in the literature.

4.1 The Concept of Genetic Algorithms

On a simple level, genetic algorithms operate on the evolution principle of “survival of the fittest”, in which the strongest characteristics or chromosomes have a higher probability of getting passed along to future generations. This forces a species to adapt its characteristics to suit the environment [81]. Adaptation takes place by crossing the chromosomes of two different organisms to produce one or more new organisms, by random mutation of one specific chromosome in one organism, and by selection for optimal chromosomes to be passed on to the next generation. There are a fixed length of chromosomes per individual, a number of individuals per population, and a number of populations evolving separately but simultaneously for each generation. Synthetic evolution in the form of genetic algorithms allows us to maintain a level of diversity in the individuals for each generation so that we are efficiently exploring the various possibilities in chromosome manifestations and combinations.

In a genetic algorithm, the string of chromosomes belonging to an organism are analogous to a set of parameters, and an individual organism is analogous to one simulated model [82]. Chromosomes are what genetically comprise an individual, and the parameter set is what is fed into the code to create a light curve model (as outlined in Chapter 3). We have a set number of individuals

to be created and tested per population, and a set number of distinct but parallel populations per generation. A fixed number of generations, to take place sequentially, comprise one full run of the genetic algorithm. The genetic algorithm also saves the more recent parts of its history with regard to parameter sets tried and how well-fitting they are, so that it can look to and learn from them when selecting future parameter sets. This will help it avoid ill-suited areas of the parameter space.

The environment is computationally analogous to the parameter space. Thus, by testing the fitness of an individual in the parameter space, we are testing how well that individual is adapted to the environment. Crossover is the primary fitness adaptation mechanism in a genetic algorithm. Two individuals are paired at random from the population, and their parameter sets are crossed to produce an offspring. Mutation, an adaptation mechanism secondary to crossover, selects one parameter at random within one individual and changes it. Fitness measures how well a given individual is suited to the environment. Our fitness parameter is the χ^2 of that individual model with the data, so that the lower the χ^2 , the better the fitness. If an individual has very good fitness, it has a higher likelihood of selection. Selection is the explicit “survival of the fittest” operator. A small percentage of the most highly fit individuals from one generation, over all the populations for that generation, become part of the starter population for the next generation. The rest of the population is filled out with crossover from these, random selection to make new individuals from parts of the parameter space that have not been sampled yet, and crossover between these new individuals and the carried-over starter individuals.

A principal concept in newer genetic algorithms is “building blocks”, or small subgroups of the full parameter set of closely-related parameters, the fitness of which can be treated like smaller sub-problems [83]. This is a particularly useful tool when dealing with a lengthy set of parameters, but is

applicable regardless of an individual’s size.

In a standard genetic algorithm, building blocks are defined as small parameter sub-groups that are used to match patterns between a large number of parameter sets. These building blocks have very good fitness, and are at the core of the selection mechanism to ensure that only the best building blocks are carried on to the next generation [83]. The smaller these blocks are, the less likely they are to be disrupted by crossover or mutation. Building blocks carve the parameter space into nearly-independent sub-spaces to be explored; this reduces a large problem into smaller, more manageable problems [84]. The parameters within a building block are also “co-adaptive”, meaning that changing each parameter individually does not improve fitness noticeably, but changing them together does improve fitness [84]. However, note that a building block can be as simple as fixing one parameter. As the genetic algorithm selects parameter sets for a new population, building block patterns that have been shown to have good fitness have a higher probability of being selected, and thus are preferentially given more new parameter sets following that pattern. In other words, these fundamental building blocks are accounting for the parameter sets of the individuals, the fitness of each individual, and the similarities among individuals within a given population [83].

We can both exchange building blocks by combining two or more sets of building blocks, and exchange parameters within the building blocks themselves. Regular crossover will cross any parameters including ones within building blocks, but building block crossover, which we can specify in Ferret (discussed in the next section), will swap building blocks.

Genetic algorithms are a robust way of fitting curves because they do not require the curves to have continuous derivatives, and they have natural methods for avoiding convergence on local minima. We are utilizing a minimizing genetic algorithm, since we want the minimum χ^2 fit with the data. Genetic algorithms are much more efficient than pure random parameter space sam-

pling due to these building blocks, while operators like mutation and crossover and selection still allow us to explore the parameter space thoroughly without wasting time on ineffective or forbidden areas.

4.2 Qubist and the Ferret Genetic Algorithm

Qubist is a global optimization platform developed by Dr. Jason D. Fiege [84] with a variety of methods for exploring the entire parameter space hunting for the global minimum. Qubist has two main powerhouse algorithms: Ferret, a genetic algorithm, and Locust, a particle swarm algorithm. We are using Ferret because its niching, dispersion, and linkage-learning (these terms will be defined in the following paragraphs) make it ideally suited to “degenerate cases” like ours [85], and there is a dearth of neutron star light curve fitting via genetic algorithm in the literature. Also, for some applications, Ferret is able to explore the parameter space more exhaustively than Locust [85].

While Ferret has the capabilities of being a “multi-objective optimizer”, we have a single objective: to minimize the χ^2 of our light curve models with the simulated data. Like a traditional genetic algorithm, Ferret uses the principles of adaptation and natural selection to find the optimal set of parameters for the objective. In addition, Ferret is capable of self-adapting to adjust the control parameters, which is unique in the genetic algorithm literature.

Ferret’s most appealing aspect is its strong linkage learning features. Linkage learning involves the code discovering for itself how the fitness of some parameters strongly affects the fitness of other parameters, and investigating which parameters are related in this way. Building blocks are composed of linked parameters, and so linkage learning discovers building blocks.

For difficult problems, such as ones with degenerate parameters and many local minima, linkage learning allows for a more thorough analysis of building blocks and how they are combined in next generations [82]. To do this, Ferret

employs not just crossover of parameters, but crossover of the building blocks as well.

A tighter linkage is a building block with a smaller “defining length”, and conversely a looser linkage has a larger “defining length”. In general, tighter linkages are better, since they are less likely to be torn apart in mutations and crossovers. Linkage learning is able to “differentiate between good and bad linkage, codings, or operators” [82], and so Ferret is able to preferentially select good linkages and therefore sturdier building blocks. Ferret learns the linkages between parameters faster than it learns what the best parameter values are [82], which sometimes requires us to explicitly slow down the convergence of values so that it will learn the parameter linkages thoroughly.

The following is paraphrased from the Qubist User’s Guide by J. D. Fiege [84].

For the majority of Ferret runs, we use 300 individuals, 15 parallel populations, and 30 generations, but for some runs we increased these values a bit for the lowest χ^2 value of the generation and the average χ^2 value of the generation to converge better. We are currently working on runs with ~ 500 generations in close collaboration with J. D. Fiege, and we anticipate publishing these results.

Ferret saves 25 generations of history; that is, all individuals and their χ^2 values for all populations within the last 25 generations. This “contain[s] all of the information necessary to analyze or resume a run, or to play a movie showing how the run progressed” [84].

Ferret’s version of the crossover function, referred to as “X-type crossover”, is an adaptation of the traditional crossover to allow for a more complete exploration of the parameter space. In Ferret the parameters are encoded as real-valued numbers, as opposed to a simplistic binary-valued bit. X-type crossover draws a line between the two mates or parents in multi-dimensional parameter space and a new offspring is selected somewhere near that line with a perturbation distance related to how far away the parents are in parameter

space. Ferret employs a dispersion tactic to fill in the crossover area between the mates, where the area of possibility for the offspring is expanded around that central line. Specifically, we use biconic dispersion, in which two statistical probability cones have a vertex on each of the mates such that the probability of selecting a new individual close to one of the mates or in between them is higher than selecting one farther away. Because of the cone shapes, there is higher probability of selecting a new individual farther away from the center dispersion line when we are midway between the two, and more chance selecting an individual close to the line when closer to one of the mates. We experimented with cylindrical, conic, and biconic dispersion, and we found the best results with biconic.

We set the probability of X-type crossovers at 100%. Since crossover is the primary method of adaptation, it is desirable to have a very high value here, so that the largest amount of parameter sets will undergo crossover. The scale of crossover is 50%, so Ferret will cross at most half of a parameter set. Within X-type crossover, Ferret also has the ability to create the offspring farther away from the other mate rather than closer to it. The probability of this anti-crossover happening is set to 0%, so crossovers will always create the offspring somewhere “in between” both mates in the parameter space. Crossover is not restricted to building blocks, so all parameters are eligible for crossover. This being said, building block crossover is considered “essential” to Ferret’s linkage learning process [84].

Due to Ferret using real-valued parameters, its mutation operator provides many more diverse possibilities than a straightforward bit flip mutator. The probability of an individual’s location in parameter space being mutated is 15%. The scale of mutation is 40%, so it will mutate a given individual 40% away within one standard deviation from the current location in parameter space [84]. Again, this is to help explore the parameter space and push Ferret out of local minima.

Ferret also has the ability to selectively mutate highly clustered individuals, to better fill out the parameter space. We have set this so that mutations happen 20% more frequently to clustered individuals than non-clustered individuals. This is designed to avoid convergence on local minima, so that over the course of the whole run Ferret is able to find the global minimum. This is very beneficial to our application, as the parameter space is riddled with local minima. Mutation is not restricted to building blocks, so all parameters are eligible. We have also allowed the probability of supermutation, a massively disruptive mutation in which “a single building block is perturbed for all non-optimal and non-elite solutions, in all populations” [84]. In very small doses this will be beneficial to us, to help avoid convergence on local minima, so we have the probability set to 1%. This is analogous to a mass-extinction of all species with mediocre adaptiveness.

Our selection pressure is 80% on overall fitness. This allows niching (clustering of solutions) to have slight priority over fitness when selecting superior individuals, and decreases the convergence rate to the benefit of exploring the parameter space more. That being said, each individual has a 100% chance of competing. The selection pressure on building blocks is 100%, so that the fitness of a building block is always a selection priority. We set 10% of the population to be devoted to exploiting the current optimal region rather than exploring. Ferret also keeps track of niching, so that it knows how clustered the individuals are in the parameter space.

Elitism makes certain that the best-fit individuals of each generation are passed to the next generation undisturbed by crossover or mutation [84]. We set the top 5% of individuals per generation to be allocated as “elite” to propagate forward.

While the populations are almost entirely independent, Ferret allows for some “immigration” between populations. We define a 1% probability that an individual will immigrate, or be randomly inserted into a different population.

We choose to keep this low so that each population will explore different areas of the parameter space independently, which gives a better chance of finding the global minimum χ^2 .

Overall, we have set Ferret’s genetic algorithm parameters such that we believe they yield good light curve fits.

The five parameters Ferret simultaneously fits are mass, radius, inclination angle, emission angle, and phase shift. We allow Ferret to explore somewhat conservative mass and radius ranges: the tested mass range is $1.0 M_{\odot}$ to $2.5 M_{\odot}$ and the tested radius range is 10 km to 16 km. Furthermore, certain areas of this mass-radius space are off limits, as they create “non-physical” neutron stars. Causality constraints do not allow for neutron stars with $\frac{M}{R} \gtrsim 0.3$ [1][4], and no physical equation of state would yield a neutron star with $\frac{M}{R} \lesssim 0.1$ (in gravitational units, where $G = c = 1$). Ferret checks these two conditions before allowing a parameter set to be used.

Inclination and emission angle each range from 0.01° to 90.0° , and normalized phase shift ranges from 0.00 to 1.00, circularly. While it may concern some that the inclination and emission angle do not extend to 0° , we find that these angles require the least amount of precision to still create a good light curve fit compared to the other parameters. On the incredibly rare occasion that we would have a neutron star with an inclination or emission angle $0^{\circ} \leq (i, \theta) < 0.01^{\circ}$, the hundredths-place precision will not pose a problem. Additionally, having $i, \theta = 0$ would not give pulsations (e.g., Figure 3.1 and Figure 3.2), and so Ferret would quickly learn to avoid these values.

4.2.1 Components of Ferret

The following are the code components that work together to comprise Ferret, listed in chronological order of when they are called within one complete Ferret run.

- *init*: Called once at the beginning of the whole Ferret run. Initializes all constants specific to *Spot* such as spin frequency, spot temperature, gray-body factor, neutron star model number, energy band intervals, angular size of the hot spot, and distance from us to the neutron star. Additionally, *init* loads the data file we will be fitting, and writes headers to a number of output files and tables that are added to as later programs run.
- *FerretSetup*: Called once at the beginning of the whole Ferret run. Declares all genetic algorithm-specific “control parameters” necessary for Ferret such as number of individuals per populations, number of populations per generation, and number of generations per run, as well as the mutation rate, crossover rate, etc., as discussed in the previous section.
- *fitness*: Called for each individual. Generates a set of parameters, checks if the mass and radius are physical, calls *spotMex* to generate the light curve, and returns the χ^2 fit of that light curve with the data.
- *spotMex*: Called in *fitness* for each individual iteration of Ferret. This works in precisely the same way as *Spot*, as defined in Chapter 3, by creating a light curve from a full set of parameters and finding that curve’s χ^2 fit with the data. *SpotMex* is rewritten as a MATLAB-readable version of *Spot* to interface properly with *fitness*.
- *outputFerret*: Called at the end of each generation. Appends the parameters of the best fit of that generation to a table, saves a plot of the best fit light curve with the data to a file so that we can visually track how well the fit is evolving, and writes an executable shell script to a file such that only the best fit from the last generation, the last run of *outputFerret*, is analyzed in secondary algorithms.
- *postProcessing*: Called once at the very end of the Ferret run. Initiates

param_degen (discussed in Chapter 5) and catalogues all Qubist-related files generated in the Ferret run so that we can revisit them if necessary for further analysis.

4.3 Other Fitting Methods

We are the first in the literature to fit neutron star X-ray light curves with a genetic algorithm to constrain the physical parameters of the neutron star. Two other fitting methods found in the literature are the Markov chain Monte Carlo method with Bayesian probabilities approach, and the Levenberg-Marquardt algorithm.

4.3.1 Markov Chain Monte Carlo + Bayesian Approach

The ubiquitous Monte Carlo method has an element of random sampling or random numbers present in the simulation or computation. This could be randomly choosing the seed for a simulation, or exploring a likelihood distribution via random sampling. The Markov chain Monte Carlo (hereafter, MCMC) method is effectively a random sampling algorithm with memory of the previous step. This correlates the sampling of the current generation with that of the previous generation [86]. There are a few different computational algorithms used to achieve the random sampling. Within the application of light curve fitting, MCMC walks around the test statistic space looking for the optimal value (e.g. exploring the χ^2 space looking for the minimum).

The basis of Bayesian statistics is in using subjective probability. It predicts the probability of a hypothesis given the observed evidence and an assumed prior probability distribution. This process can be done iteratively — with more evidence, the probability estimate is updated. In light curve fitting, this is looking at how the test statistic improves when moving in one direction versus another in the parameter space, and then using this knowledge to give prefer-

ence to some subsets of the parameter space over others. By using Bayesian statistics to correlate the random samples of the current generation with the previous generation, MCMC is an effective intelligent random walk algorithm.

There is a “burn-in” period in which the algorithm explores the test statistic space to get an idea of its topology. After the burn-in period it starts to look for the minimum. Technically, the Bayesian probability of a marginal likelihood (probability of an event happening, i.e., a parameter having a specific value) is a multidimensional integral, and the integrand is what the MCMC is helping to integrate. At the completion of the run, the most probable parameter set is declared to be the best fit.

While the MCMC + Bayesian approach is very computationally expensive, this method is well-suited to exploring and optimizing poorly-behaved parameter spaces.

The MCMC + Bayesian approach is exploding in popularity in the literature, for neutron star light curve fitting as well as light curve fitting from transiting exoplanets, gamma ray and radio emission from standard millisecond pulsars, and gravitational lensing. Analyses of note are Bhattacharyya et al. 2005 [45], Steiner et al. 2010 [87], and Lo et al. 2013 [39].

In Bhattacharyya et al. 2005 [45] the authors re-create light curves to match those observed from XTE J1814–338. They do not go into much detail in their use of MCMC, but they appear to be using it in conjunction with a grid method to step through the parameter space to find areas of low χ^2 . There is no mention of a Bayesian approach in their analysis.

In Steiner et al. 2010 [87] the authors determine possible radii for a $1.4 M_{\odot}$ neutron star to constrain dense matter equations of state. They use MCMC + Bayesian to determine the “probability densities of quantities of interest”, or the likelihood that a specific parameter value would be a good fit.

In Lo et al. 2013 [39] the authors create model light curves and then fit them to constrain the original parameters, much like we show in this thesis.

They use a MCMC + Bayesian algorithm to explore the parameter space and then compute the Bayesian probability of each parameter set, where the most probable is considered the best fit with their simulated data.

In Wang et al. 2013 [88] the authors fit multi-band optical spectra of the SAX J1808 binary system during an X-ray burst and in quiescence (non-outburst) to obtain constraints on parameters of the neutron star and companion. They use a MCMC + Bayesian approach to select the parameters of the fitting model and test the strength of each fit.

4.3.2 Levenberg-Marquardt Method

The Levenberg-Marquardt method is a first-order optimization algorithm that implements iterative least-squares curve fitting, or trying to minimize the sum of the square of the deviation between fit value and data. This has a tendency to get stuck on and converge to the first local minimum it encounters. This method requires that the parameters are selected first, and then the fit is found.

While the Levenberg-Marquardt method is not advisable for optimizing challenging parameter spaces such as those we encounter, it is pre-made and readily available in commercial software packages like MathCAD. This ease of use makes it a useful tool to confirm a global minimum as determined by another method like genetic algorithms or MCMC+Bayesian.

Analyses using the Levenberg-Marquardt method include those by D. A. Leahy in [32][26][89]. In these papers the authors use the algorithm to step through parameter sets with a variable step size looking for minima in the least squares fit. To bypass local minima they can select the starting parameter set for later in the iteration sequence.

In the next chapter we discuss a secondary algorithm which analyzes the parameter degeneracies of optimal parameter set from Ferret.

Chapter 5

Understanding Parameter Degeneracies

Parameter degeneracy is the conflation of effects from different combinations of parameters. In application to our research with light curves, it means the effects that one parameter has on a light curve can be easily mimicked by the effects of one or more other parameters. This is a major hurdle to using light curves to ascertain the parameters of the neutron star that created it.

Extracting parameters from a light curve is not a simple procedure. Once a best fit parameter set is obtained, we want to be confident in the accuracy of these results by knowing the extent of equivalences or degeneracies present in the parameter space. Using the Ferret genetic algorithm to find the minimum χ^2 parameter set, we set out to determine the robustness of this fit by searching the parameter space for other solutions as good as Ferret's best fit. In this chapter we motivate and define an algorithm for finding comparable solutions.

The first part of this chapter explains how we characterize a light curve and how this will be used to explore parameter degeneracies. The second part of this chapter describes a secondary algorithm written to explore the parameter degeneracies and alternative fits for one of Ferret's optimal solutions. The third part of this chapter discusses the parameter degeneracies we expect in

our parameter sets.

5.1 Characterizing Light Curves

The two attributes we use to characterize a light curve are normalized pulse amplitude and asymmetry due to Doppler boosting, as in Figure 5.1. The fundamental idea for the secondary algorithm described in the next section is that if two light curve models have the same pulse amplitude and Doppler boost, they will have very similar (if not the same) χ^2 fits with a given data set.

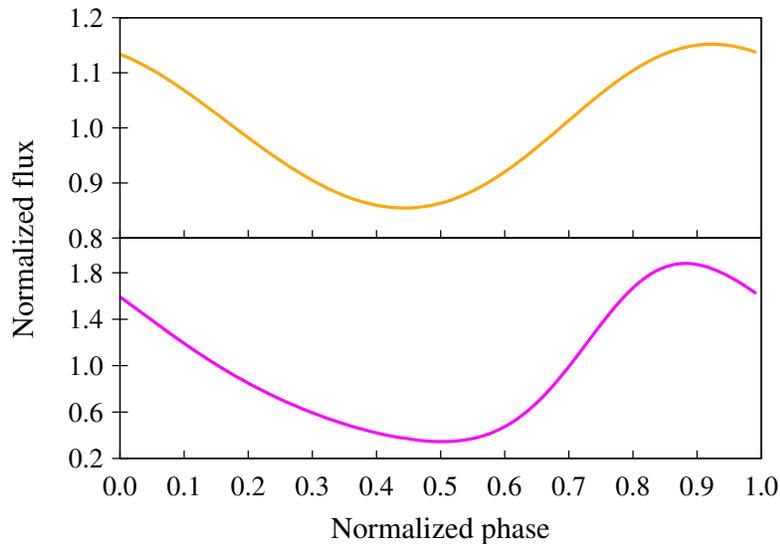


Figure 5.1: Light curves from two different neutron star parameter sets, demonstrating different pulse amplitudes and asymmetries due to Doppler boosting. The top light curve has a low pulse amplitude and a negligible Doppler boost asymmetry. The bottom light curve has a high pulse amplitude and a pronounced asymmetry from strong Doppler boosting. The two light curves are in phase, as indicated by the shared x -axis, but have different normalized flux scales on the y -axes. These light curves are shown in the low energy band, 2–3 keV.

In the secondary algorithm we will be using approximations for the pulse

amplitude and Doppler boost. We establish here that this approximation is not very close to the true value, but changes in the approximation value do translate to similar changes in the true value, and thus it is a useful proxy for comparing two light curves.

The true pulse fraction of a light curve is given by

$$A = \frac{F_{\max} - F_{\min}}{F_{\max} + F_{\min}}. \quad (5.1.1)$$

An approximation prevalent in the literature is the Beloborodov pulse amplitude approximation [56],

$$A = \frac{\left(1 - \frac{2M}{R}\right) \sin i \sin \theta}{\frac{2M}{R} + \left(1 - \frac{2M}{R}\right) \cos i \cos \theta}. \quad (5.1.2)$$

This approximation does not take into account effects of Doppler boosting on the pulse amplitude. We use a first-order approximation for Doppler boosting,

$$B = \frac{R\Omega \sin i \sin \theta}{\sqrt{1 - \frac{2M}{R}}}, \quad (5.1.3)$$

where Ω is the angular spin frequency of the neutron star.

Table 5.1 contains the pulse amplitudes as approximated by equation (5.1.2) compared to the true pulse amplitudes as computed by equation (5.1.1). While the approximate pulse amplitude does not match any of the four true pulse amplitudes we compute, there is only a $\sim 10\%$ discrepancy between the low energy band true pulse amplitude and the approximate pulse amplitude. We are further reassured that the difference between the pulse amplitudes for two very different physical systems is much larger than this discrepancy.

	1 st curve	2 nd curve
Approx. PA	0.132	0.624
True PA, bolo	0.160	0.730
True PA, mono	0.146	0.682
True PA, low	0.148	0.690
True PA, high	0.169	0.751

Table 5.1: Approximated and exact pulse amplitudes for two light curves. “1st curve” indicates the top light curve, and “2nd curve” indicates the bottom light curve in Figure 5.1. Pulse amplitude is abbreviated as “PA” here. The first row is the approximated pulse amplitude as calculated by equation (5.1.2); the second row is the true pulse amplitude of the bolometric flux; the third row is the true pulse amplitude of the monochromatic flux at 2 keV; the fourth and fifth rows are the true pulse amplitudes of energy band flux in the 2–3 keV range and 5–6 keV range, respectively. All true pulse amplitudes are calculated with equation (5.1.1). We notice that the approximated pulse amplitudes are closest to but a bit lower than the monochromatic flux pulse amplitudes.

5.2 Degeneracy Contour Code

Using the approximate pulse amplitude and an approximation of the Doppler boosting in a light curve, we explain the development of a code, *param_degen*, which solves for other parameter sets similar to the best fit light curve found by Ferret. The motivation for developing this code is to find parameter sets with similar χ^2 values as Ferret’s best fitting parameter set, since the χ^2 contour regions are not straightforward in the mass-radius plane. We want to ensure that other solutions as good as the best fit are also accounted for.

5.2.1 Solving for Inclination and Emission Angles

Given a parameter set $\{M, R, i, \theta\}$ that has A as the Beloborodov approximation for the pulse amplitude and B as an approximation of the Doppler boost in the curve, we want to find what range of values of M, R will have a similar value of A and B . First we choose a value of M and R by looping through an evenly spaced “grid” of masses and radii that covers the allowed mass-radius

space. Next we find possible values of i and θ while fixing A and B . If a solution for i and θ exists for a given M and R , then we run *Spot* with the values to create a light curve and obtain its χ^2 fit with the simulation data.

We have the Beloborodov approximation for the pulse amplitude, given by equation (5.1.2), and a first-order approximation of the Doppler boosting given by equation (5.1.3). Both A and B are known for a given $\{M, R, i, \theta\}$.

We loop through a range of M and R values, denoted m and r . We define

$$\zeta = \frac{B \sqrt{1 - \frac{2m}{r}}}{r \Omega} \quad (5.2.1)$$

such that

$$\zeta = \sin i \sin \theta, \quad (5.2.2)$$

so

$$A = \frac{\left(1 - \frac{2m}{r}\right) \zeta}{\frac{2m}{r} + \left(1 - \frac{2m}{r}\right) \cos i \cos \theta}. \quad (5.2.3)$$

Now we solve for the two unknowns, i and θ , from equations (5.2.2) and (5.2.3).

For equation (5.2.2) we must have $0 < \zeta < 1$, and so

$$\cos i \cos \theta = \frac{\sin i \sin \theta}{A} - \frac{2m}{r} \frac{1}{\left(1 - \frac{2m}{r}\right)}. \quad (5.2.4)$$

We define κ such that

$$\kappa = \frac{\zeta}{A} - \frac{2m}{r} \frac{1}{\left(1 - \frac{2m}{r}\right)}, \quad (5.2.5)$$

so

$$\kappa = \cos i \cos \theta \quad (5.2.6)$$

with the rule $0 < \kappa < 1$, so

$$\kappa^2 = \cos^2 i \cos^2 \theta. \quad (5.2.7)$$

Using trigonometric identities and equation (5.2.2), we get equation (5.2.7) in

terms of one variable and manipulate it to take the form of a quadratic:

$$\kappa^2 = (1 - \sin^2 i)(1 - \sin^2 \theta) \quad (5.2.8)$$

$$\kappa^2 = (1 - \sin^2 i) \left(1 - \frac{\zeta^2}{\sin^2 i}\right) \quad (5.2.9)$$

$$\kappa^2 \sin^2 i = (1 - \sin^2 i)(\sin^2 i - \zeta^2) \quad (5.2.10)$$

$$\kappa^2 \sin^2 i = -\sin^4 i + \sin^2 i + \zeta^2 \sin^2 i - \zeta^2 \quad (5.2.11)$$

$$0 = \sin^4 i + (\kappa^2 - \zeta^2 - 1) \sin^2 i + \zeta^2. \quad (5.2.12)$$

This is now a quadratic in $\sin^2 i$. Using the quadratic formula, we find

$$\sin^2 i = \frac{-(\kappa^2 - \zeta^2 - 1) \pm \frac{1}{2} \Delta^{\frac{1}{2}}}{2}, \quad (5.2.13)$$

where

$$\Delta = (\kappa^2 - \zeta^2 - 1)^2 - 4\zeta^2, \quad (5.2.14)$$

and we require $\Delta \geq 0$, $\sin^2 i \geq 0$, and $0 \leq \sin i \leq 1$. Note that equation (5.2.13) has two solutions stemming from the \pm sign. In Chapter 6 we will refer to solutions that take the $+$ as “positive solutions” and ones that take the $-$ as “negative solutions”.

If a solution for $\sin i$ exists, then, by equation (5.2.2),

$$\sin \theta = \frac{\zeta}{\sin i}, \quad (5.2.15)$$

with the requirement $0 \leq \sin \theta \leq 1$, $\zeta < \sin i < 1$, and $\zeta < \sin \theta < 1$. Thus if a solution for $\sin \theta$ exists, then there exists a solution $\{m, r, i, \theta\}$ with a given A and B .

In this code, we allow for a $\pm 2\%$ tolerance in A and $\pm 3\%$ tolerance in B for low noise cases, and $\pm 8\%$ and $\pm 6\%$ tolerance in A and B , respectively, for high noise cases. These values are decided from comparing A and B values between

the best fit and ideal parameter sets, as noted in Table 6.3. There are seven divisions in both, for a total of 49 AB combinations, with one of them having the precise value of A and B as Ferret’s best fit solution. We loop through these in a nested loop that sits inside the mass and radius nested loops.

5.2.2 Components of the Degeneracy Contour Code

Both of the following programs are called in *postProcessing* at the end of a Ferret run. Here we outline the function of each program, in call sequence:

- *param_degen*: Generates a set of parameters with the same pulse amplitude and Doppler boosting as Ferret’s best fit parameter set using the algorithm outlined in the previous section. It then computes the χ^2 fit of the light curve for each generated parameter set with the data for that case and saves this information to a table, to be plotted by *chi2contours*. We refer to *param_degen* as the secondary algorithm to Ferret.
- *chi2contours*: Sorts all parameter sets created in *param_degen* by χ^2 into σ deviations from Ferret’s best fit, from 1σ to 5σ . It then plots these in the 2-dimensional mass-radius space with different colours for each σ , so that each point in Figure 6.6 is one parameter set.

We combine these results with those from Ferret to give the solutions and constraints in Chapter 6.

5.3 Parameter Degeneracies in Our Research

Parameter degeneracies are an expected obstacle in fitting neutron star light curves. Here we explain the origin of common parameter degeneracies.

For all system geometries examined, increasing the inclination angle or the emission angle correspondingly increases the observed pulse amplitude, as evident in Figures 3.1 and 3.2. This is caused by decreasing the ability to see

the hotspot for farther around the neutron star, or seeing a smaller area of the hotspot for a larger fraction of the phase, which in turn increases the observable pulse modulation. Equation (5.1.2) verifies our observations: as $(i, \theta) \rightarrow 90^\circ$, we get $\sin(i, \theta) \rightarrow 1$ and $\cos(i, \theta) \rightarrow 0$, which manifests as a larger pulse amplitude.

Some inclination and emission angle combinations yield such an enormous pulse amplitude that an eclipse occurs, in which the hotspot is not visible for a portion of the phase. Eclipses are theoretically replicable but are not evident in most observations of thermonuclear burst oscillations.

When both inclination and emission angle are high, stronger Doppler effects are observed as well as a large pulse amplitude. The inclination and emission angles are easily swapped by Ferret, as evident by the i - θ symmetry in equations (5.1.2) and (5.1.3). However there are subtle differences in the fits of one compared to another that are not caught by our approximation, so in practice they are not entirely interchangeable.

Increasing compactness (that is, increasing the mass for a fixed radius, or decreasing the radius for a fixed mass) decreases pulse amplitude. With a more compact neutron star, the gravitational light bending is stronger, so the observer is able to see more of the total surface area of the neutron star. This lets the observer see a larger area of the hotspot for a larger fraction of one rotation, which decreases the overall variability in brightness, which decreases the observed pulse amplitude. In looking at equation (5.1.2), increasing $\frac{M}{R}$ will increase the denominator more than the numerator, which corresponds to a decrease in pulse amplitude.

Doppler effects also play a role in parameter degeneracies, namely that it is more difficult to have parameter degeneracies with a strongly Doppler boosted light curve. Strong Doppler effects are accomplished by increasing the relative velocity of the hotspot to the observer. We can achieve this by giving the neutron star a larger radius and have the spot at a higher emission angle,

having both a high inclination angle and a high emission angle, or increasing the spin frequency of the neutron star. This is evident in equation (5.1.3): as R or Ω increases we will see an increase in B , or as $(i, \theta) \rightarrow 90^\circ$ we have $\sin(i, \theta) \rightarrow 1$, which also boosts B . By combining conditions or exaggerating a condition, we see a stronger Doppler boost and thus a more amplified asymmetry in the light curve.

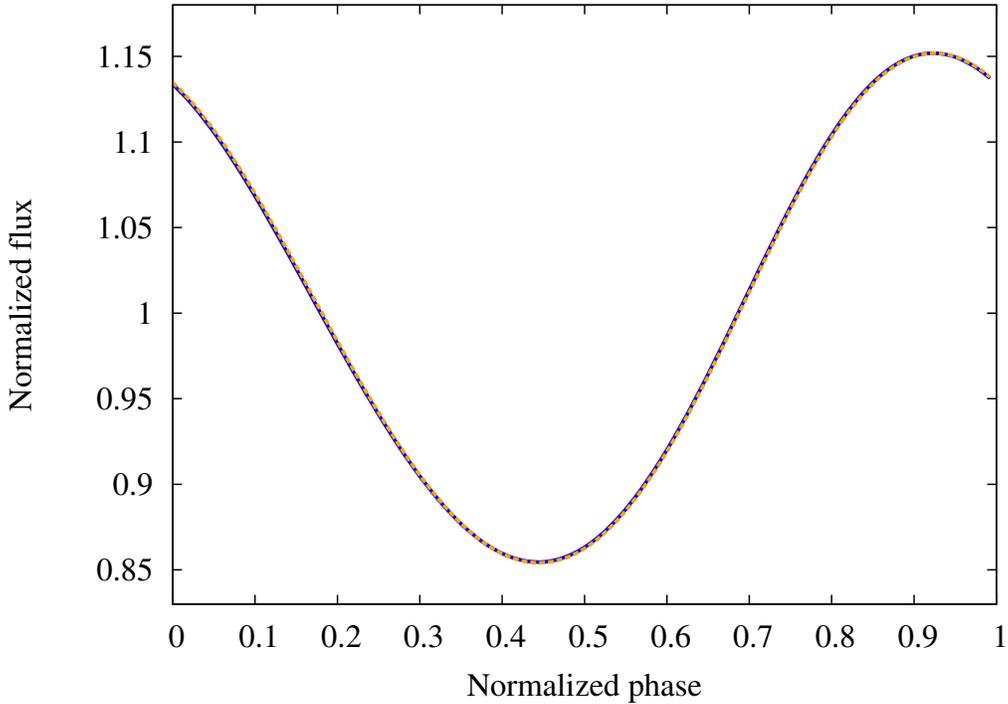


Figure 5.2: Simulated light curves from two different neutron stars. Light curve 1 (blue solid) has $M = 1.60 M_\odot$, $R = 12.0$ km, $i = 10.0^\circ$, $\theta = 60.0^\circ$, and no phase shift. Light curve 2 (orange dotted) has $M = 1.55 M_\odot$, $R = 10.8$ km, $i = 61.0^\circ$, $\theta = 10.6^\circ$, phase shift = 0.999. All other parameters are identical. Shown in the low energy band, 2–3 keV. The two light curves are indistinguishable. See text for discussion.

Figure 5.2 elucidates the concept of parameter degeneracy. It shows two simulated light curves from two distinct neutron star parameter sets. The neutron star that created light curve 1 is less compact than the neutron star that created light curve 2, and it has a low inclination and high emission angle. The inclination and emission angle values for neutron star 1 are approximately

flipped for neutron star 2. However, these light curves appear to be identical. We can infer that the compactness for neutron star 2, which would make the light curve less modulated, is balanced by its flipped and slightly higher values of inclination and emission angle, which would make the light curve more modulated. This compensation of the pulse amplitude effects yields carbon copy light curves.

We use the *param_degen* code to understand the results from and limitations of fitting light curves with the Ferret genetic algorithm in the next chapter.

Chapter 6

Results

In this chapter we discuss our light curve fitting results from the Ferret genetic algorithm and the constraints on those fits given by the secondary algorithm *param_degen*.

We use Ferret to fit light curve models to a data set to determine the best parameter set to describe the data, and therefore the likely parameters of the system that created the light curve. We simulate data as perturbations from an ideal curve, so that we know the true parameters used to make each data set, but we envision using the information we find about Ferret’s accuracy (from an algorithmic standpoint and a parameter degeneracy standpoint) to fit real light curves and help place constraints on masses and radii of neutron stars. Here we have ten different cases illustrating different system geometries.

We fix the spin frequency, graybody emission toggle, and neutron star shape model, noted for each case in the following subsections. Additionally, we have a fixed distance of 6 kpc, a hotspot angular radius of 6° , a 1×1 hotspot mesh, and a hotspot temperature of 2 keV. For all but one case we select a neutron star spin frequency of 581 Hz, the measured spin frequency of the Type I X-ray burst source 4U 1636–536. In this chapter we have 59 degrees of freedom (d.o.f.), since we are fitting the low energy band and high energy band simultaneously: this gives $2 \times 32 - 5 = 59$ d.o.f., for the 32 phase bins per energy band, subtracted

by the five fitting parameters.

The ideal curves in Figures 6.1, 6.2, 6.8, 6.10, 6.12, 6.14, 6.16, 6.18, 6.20, 6.22, and 6.24 were computed using the *Spot* code. The simulated data points in Figures 6.1, 6.2, 6.8, 6.12, 6.14, 6.16, 6.18, 6.20, and 6.22 were computed from their respective ideal curves using Poisson statistics by Dr. Denis A. Leahy. These error bar sizes and noise levels are similar to what we would expect to get from a very sensitive instrument such as LOFT. Data points in Figures 6.10 and 6.24 were computed using high noise statistics, based on observations of 4U 1636–536 from RXTE.

The fit curves in Figures 6.1, 6.2, 6.8, 6.10, 6.12, 6.14, 6.16, 6.18, 6.20, 6.22, 6.24 are the light curves of the parameters that Ferret determined to be the best fits to the data points. We ran Ferret twice for the high noise data sets and thrice for the low noise data sets, and for each case we picked the best of the best-fit runs to use as the simulation parameters in the figures.

To compute the accuracy of the best-fit mass and radius parameters, we use the following equations for percent discrepancy:

$$\%_{\text{err}_M} = \frac{|M_{\text{fit}} - M_{\text{ideal}}|}{M_{\text{fit}}}, \quad (6.0.1)$$

$$\%_{\text{err}_R} = \frac{|R_{\text{fit}} - R_{\text{ideal}}|}{R_{\text{fit}}}. \quad (6.0.2)$$

Note that we use the fit value in the denominator, instead of the ideal value. This is because when analyzing observed data, we will only know the fit values and not the ideal or true values, and so we will want to give our accuracy as percent discrepancy from the fit values.

6.1 Case 1: The Fiducial Case

The Case 1 light curve is characterized by a symmetric sinusoidal shape and a low pulse amplitude. For this case we did nine runs of Ferret, and chose the

lowest χ^2 fit for the subsequent plots and discussion. We did a number of runs for this case to understand the accuracy and errors in Ferret’s convergence for the rest of the test cases, to determine how much fluctuation was possible in the best-fit values. The optimal parameter set for each run is listed in Table 6.1. Of the nine runs, the largest percent discrepancy in mass was 10.0% and the largest percent discrepancy in radius was 14.4%. We did not calculate percent discrepancy for inclination or emission angle since they are not crucial parameters. We deduce that for symmetric curves with low pulse amplitude, it is likely that Ferret will fit the mass and radius within $\sim 10\%$ and $\sim 15\%$ of the correct value, respectively.

In Figures 6.1 and 6.2 we can see that the ideal and fit curves are incredibly close, and both replicate the data very well. The χ^2 of the fit with the data is lower than the χ^2 of the ideal with the data. However, due to parameter degeneracies, the parameter sets are notably different (see Table 6.1). We notice that inclination angle and emission angle fits seem to be switched, and for all the test cases (see Table 6.3) the two appear to be interchangeable. In examining the parameters in Case 1 Fit 9, an inclination angle fit of 60.97° is not remotely close to the ideal 10° , but it is an excellent fit for the ideal emission angle of 60° . Likewise, the emission angle fit of 10.59° is a horrendous fit for the ideal 60° , but very close to the ideal inclination angle of 10° . The fit mass is 3.1% from the ideal, and the fit radius is off by 9.8%.

Since we can see that the fit curves are very close to the data points in both the low and high energy bands, we will only show the low energy band light curves for comparison with the next cases.

6.1.1 Ferret

Figure 6.3 shows the mass and radius values of every individual from Ferret that is within 5σ of the best fit for the run. There are no defined σ contours in the mass-radius plane, likely because this plot only shows χ^2 as a function

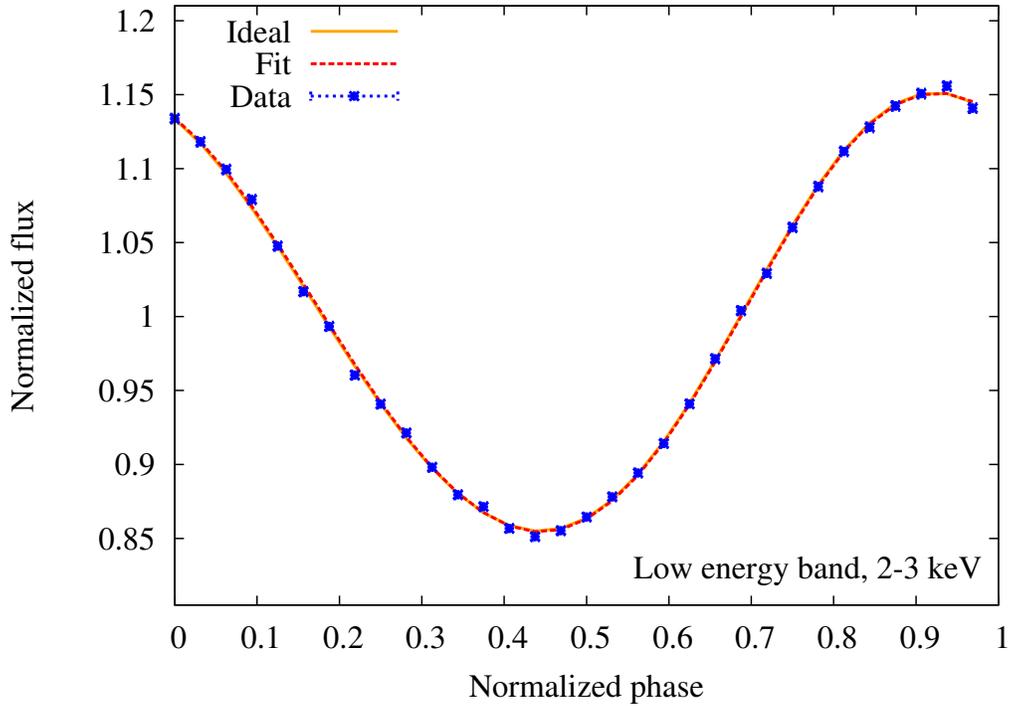


Figure 6.1: Case 1: ideal curve (orange solid), Ferret’s best-fit curve (red dotted), and simulated data (blue) for a spherical neutron star with isotropic emission, rotating at 581 Hz. Ideal parameters: $M = 1.60 M_{\odot}$, $R = 12$ km, $i = 10^{\circ}$, $\theta = 60^{\circ}$, phase shift = 0, $\chi^2 = 58.88$. Fit parameters: $M = 1.55 M_{\odot}$, $R = 10.82$ km, $i = 60.97^{\circ}$, $\theta = 10.59^{\circ}$, phase shift = 0.9991, $\chi^2 = 55.43$. Plotted in the low energy band, 2–3 keV. For reference, the ideal curve has a gravitational redshift of $1 + z = 1.28448$.

of two parameters, out of five total parameters. There appear to be three 1σ deep local minima, with one of them being the global minimum as determined by Ferret.

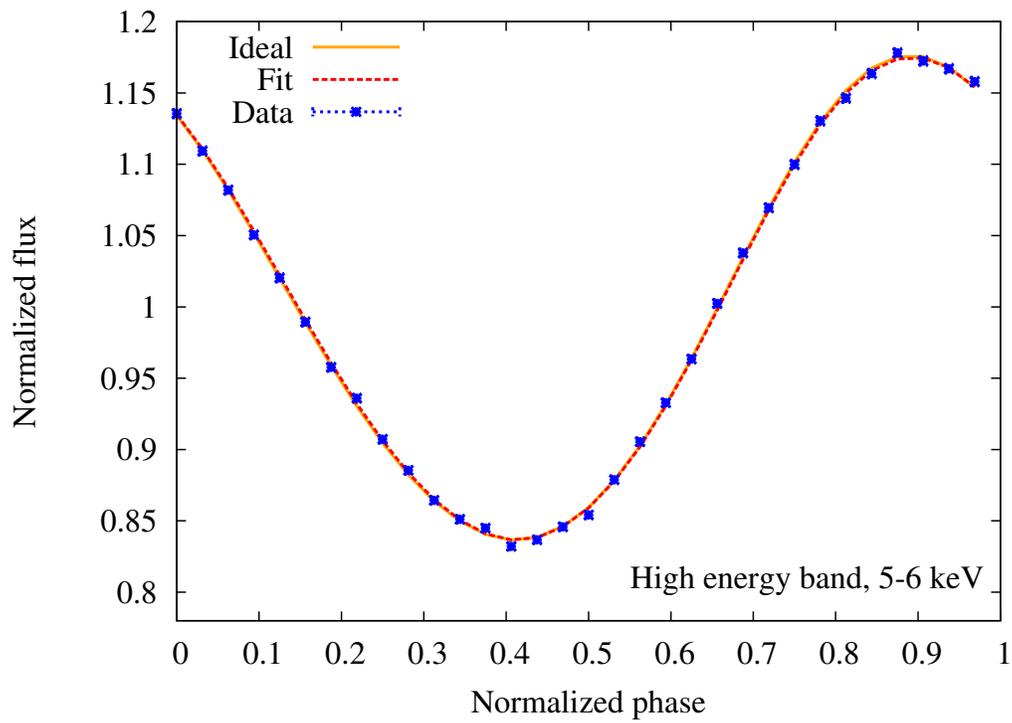


Figure 6.2: Case 1: ideal curve (orange solid), Ferret’s best-fit curve (red dotted), and simulated data (blue) with the same parameters as in Figure 6.1. Plotted in the high energy band, 5–6 keV.

	χ^2	M (M_\odot)	% err _M	R (km)	% err _R	i ($^\circ$)	θ ($^\circ$)	$phase$	Num. σ
Ideal	58.88	1.60	-	12.00	-	10.00	60.00	0.00	-
Fit 1	57.26	1.64	2.5%	12.67	5.6%	64.47	8.910	0.9957	1 σ
Fit 2	56.67	1.65	3.1%	12.88	7.3%	8.768	64.52	0.9959	1 σ
Fit 3	56.11	1.63	1.9%	11.81	1.6%	64.66	9.424	0.9964	2 σ
Fit 4	58.46	1.66	3.8%	13.73	14.4%	8.595	62.84	0.9972	1 σ
Fit 5	55.55	1.55	3.1%	11.24	6.3%	10.12	61.49	0.9978	2 σ
Fit 6	57.33	1.59	0.6%	11.21	6.6%	60.71	10.48	0.0010	1 σ
Fit 7	57.42	1.64	2.5%	10.72	10.7%	74.11	8.917	0.9914	1 σ
Fit 8	58.18	1.44	10.0%	11.46	4.5%	12.01	50.44	0.0028	1 σ
Fit 9	55.43	1.55	3.1%	10.82	9.8%	60.97	10.59	0.9991	2 σ

Table 6.1: Results for Case 1. Each line is from one full run of Ferret, for nine total runs. Columns 4 and 6 are the percent discrepancy of the mass and radius fit values, respectively. The largest percent discrepancy in mass is 10.0% and the largest percent discrepancy in radius is 14.4%. The last column, “numerical σ ”, is within what σ of Ferret’s fit the ideal solution falls in, as computed using $\Delta\chi^2$ for two free parameters; $\Delta\chi^2 = 2.295$ for 1 σ , $\Delta\chi^2 = 6.180$ for 2 σ , $\Delta\chi^2 = 11.83$ for 3 σ , $\Delta\chi^2 = 19.34$ for 4 σ , and $\Delta\chi^2 = 28.74$ for 5 σ . All of the fits encompass the ideal by at most 2 σ . We see that Fit 9 has the lowest χ^2 , and so we use this parameter set for the plots in this section. Notice that all of the fit χ^2 values are lower than the ideal χ^2 . There are 59 degrees of freedom.

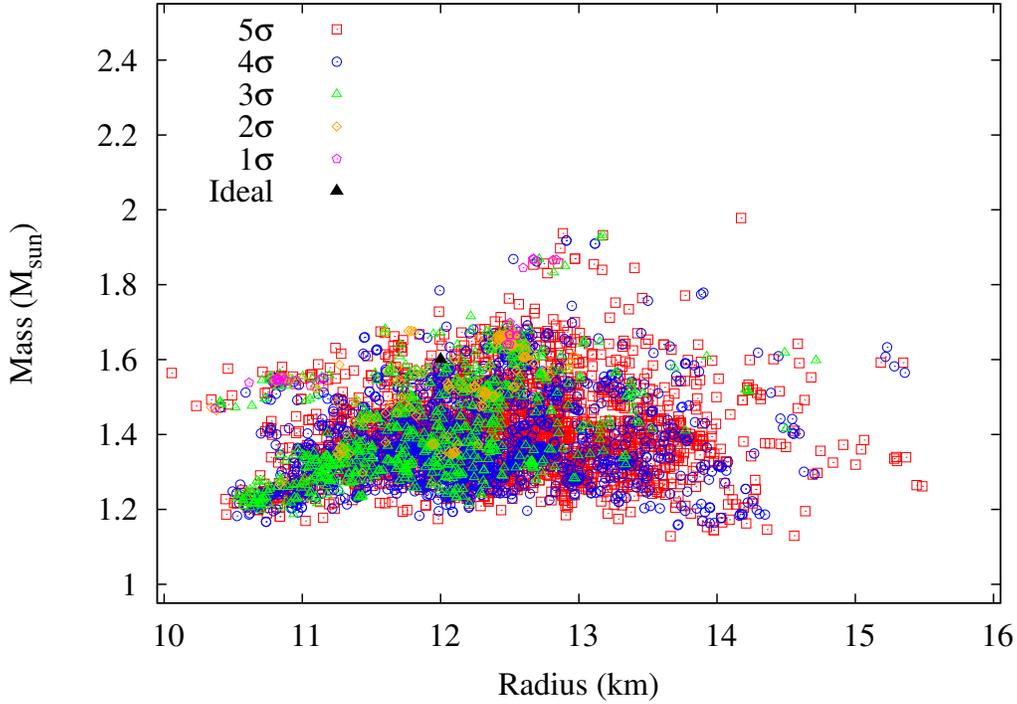


Figure 6.3: Case 1: The mass and radius values from all individual parameter sets within 5σ of the optimal parameter set. Each point represents one individual model that was run by Ferret. The σ contour levels were calculated from Fit 9, $\chi^2 = 55.43$, which had $M = 1.54 M_\odot$ and $R = 10.82$ km. The χ^2 cut-offs are 57.73 for 1σ , 61.61 for 2σ , 67.26 for 3σ , 74.77 for 4σ , and 84.17 for 5σ . The ideal mass and radius used to create the simulated data is noted with a black triangle at $1.6 M_\odot$ and 12 km.

Additionally, if we examine one very small area of the plot, the mass and radius values of different points in that area are not that different, but the points have different χ^2 values and σ levels away from the minimum. This demonstrates that the inclination angle, emission angle, and phase shift values clearly have an effect on how well the light curve fits the data.

To visualize how these individuals sit in the parameter space, we plot the minimum χ^2 surface versus mass and radius. This helps us understand the topology of the parameter space and where the best-fit solutions fall.

To make a surface plot of the χ^2 values against mass and radius (Figure 6.4), we “grid” the data to create a surface plot within Gnuplot. Each data

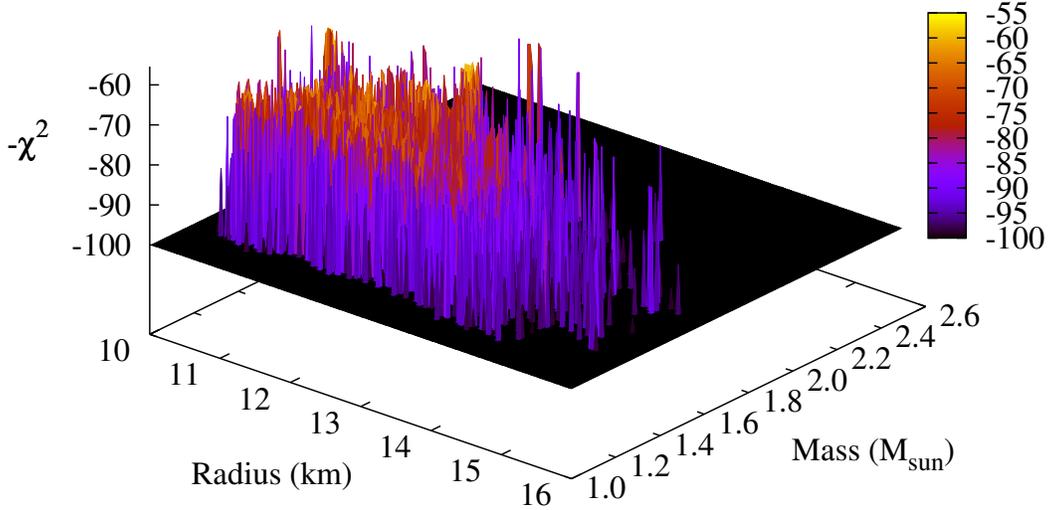


Figure 6.4: Case 1: Surface plot of the χ^2 between the simulated data and the parameter set with each respective M and R value for all individuals run by Ferret for Fit 9. We plot models with $\chi^2 \leq 100$, since we aim to minimize χ^2 . Since Gnuplot only plots maximal surfaces, we plot $-\chi^2$ on the z -axis, so what appear to be maxima here are actually minima. Due to the way that the data is binned to make it fit the surface plot requirements, the plot is a little spikier than expected, but we do still expect many local minima.

point takes on the middle value of the mass bin, middle value of the radius bin, and the lowest χ^2 value of the individuals in that MR bin. We are using 200 divisions in both mass and radius, for 4000 MR bins total; we want a smoother surface, for which larger bins would make sense, but we also do not want to gloss over the local minima and bumpy parameter space landscape, which require smaller bins for accurate visualization. The plot shows many local minima, a few more significant local minima, and one global minimum containing Ferret’s best-fit solution.

Then we colour-mapped Figure 6.4 onto the mass-radius plane to get Figure 6.5. As with Figure 6.4, we are mapping $-\chi^2$. This makes the minima (now

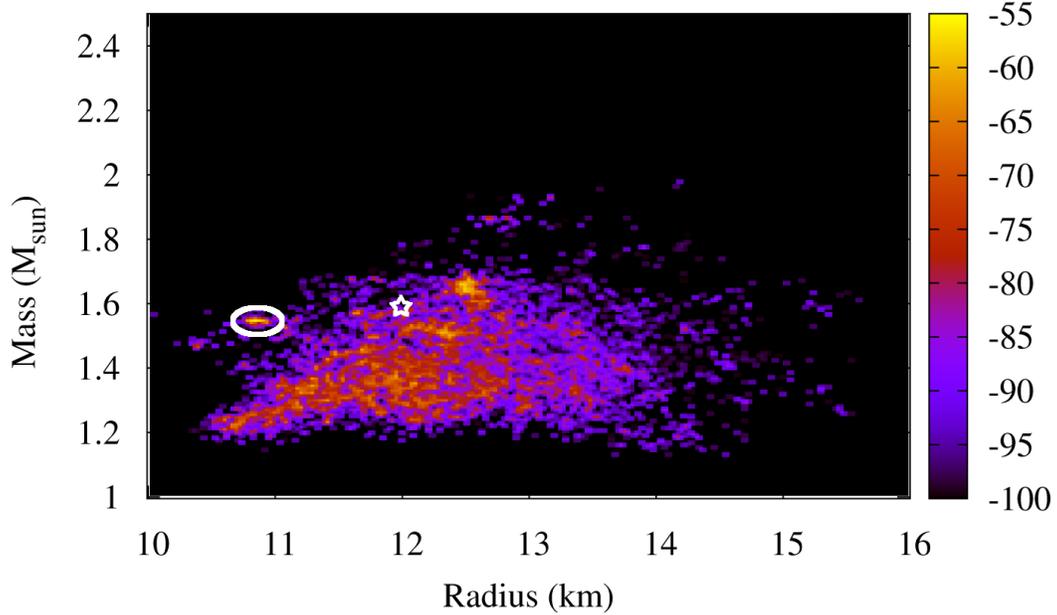


Figure 6.5: Case 1: The surface plot of Figure 6.4 colour-mapped onto the $M R$ plane. We plot $-\chi^2$ values with the colour map because Gnuplot surfaces are maximize-only. The ideal parameters are $M = 1.6 M_\odot$ and $R = 12$ km with $\chi^2 = 58.88$, marked with a star, and the best-fit parameters are $M = 1.54 M_\odot$ and $R = 10.82$ km with $\chi^2 = 55.43$, within the circled minimum.

maxima) easier to see, and so the yellow colour represents $\chi^2 \sim 55$. The local minima are much more apparent on this plot than Figure 6.4. There are a few minima regions, with two standing out as the lowest local minima. The small thin elliptical minima region on the left side of the plot contains Ferret’s best-fit solution. Since this data is a coarser version of that in Figure 6.3, we are not able to see the finer features of the χ^2 surface such as more local minima, but the two plots have the same overall qualities.

In looking at Figures 6.3, 6.4 and 6.5, the most amount of information in an easily presentable manner is in the Ferret χ^2 contours of Figure 6.3. Therefore we elect to show only the Ferret χ^2 contour figures for Cases 2–10.

6.1.2 Parameter Degeneracy

We use the parameter sets identified in *param_degen* and their respective χ^2 fits to map out the extent of the degeneracies in parameter space. This is to detect other areas of the parameter space that fit the simulated data well that may have been missed by Ferret. We are using Ferret’s best-fit solution, Fit 9, for Case 1, since it had the lowest χ^2 of all nine runs completed for Case 1.

The χ^2 contour plot in Figure 6.6 is created in *chi2contours*, as explained in section 5.2. Each point in Figure 6.6 is one parameter set run by *param_degen*. These points are then sorted by $\Delta\chi^2$ region and plotted in the mass-radius plane. The 5σ parameter sets are plotted first, followed by the 4σ , 3σ , 2σ , and 1σ points in superposition. This is so that if a (M, R) point has one solution that is 2σ but another solution that is 1σ , it will appear 1σ on the plot, since the other parameters and characteristics are less crucial. We care more about whether a given (M, R) is able to produce an optimal parameter set and less about what the other values of that parameter set are. The $\Delta\chi^2$ for σ levels for two free parameters, mass and radius, are as follows: $\Delta\chi^2 = 2.295$ for 1σ , $\Delta\chi^2 = 6.180$ for 2σ , $\Delta\chi^2 = 11.83$ for 3σ , $\Delta\chi^2 = 19.34$ for 4σ , and $\Delta\chi^2 = 28.74$ for 5σ .

Furthermore, we plot the ideal solution, used to create the simulated data set, on the χ^2 contour plots, to get an idea of how well Ferret and *param_degen* can predict the true mass and radius of a given data set. Compared to Figure 6.3, Figure 6.6 shows that the χ^2 space from *param_degen* is much smoother and more gradual than that discovered by Ferret.

In Figure 6.6 some of the points appear to have “stacked” or overlapping σ levels. This is due to having a small tolerance in possible pulse amplitude and Doppler boost solutions (equations (5.1.2) and (5.1.3)) for a given mass and radius, and to having positive and negative solutions to equation (5.2.13), which could give somewhat different χ^2 values. This suggests that the emission and inclination angles are not as interchangeable as originally thought (as indicated

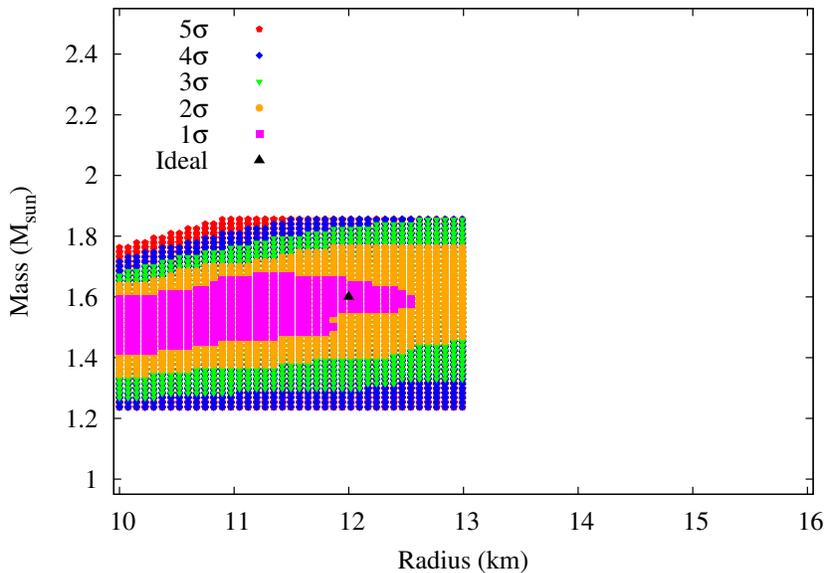


Figure 6.6: Case 1: A map of the χ^2 space up to 5σ away from the best-fit χ^2 for Fit 9, contour-mapped onto the MR plane. Fit 9 has $M = 1.54 M_\odot$, $R = 10.82$ km, and $\chi^2 = 55.43$. The ideal parameter set, which was used to create the simulated data for Case 1, is plotted with a black triangle. It has $M = 1.6 M_\odot$, $R = 12$ km, and $\chi^2 = 71.99$. We see that the ideal is within the 1σ region of Ferret’s best-fit. The χ^2 cut-offs are 57.73 for 1σ , 61.61 for 2σ , 67.26 for 3σ , 74.77 for 4σ , and 84.17 for 5σ . We notice that some of the points on the χ^2 contour boundaries overlap; see discussion in text. The abrupt cut-off in mass and radius is due to setting the mass and radius range to be $\pm 20\%$ of the best-fit values for analysis with *param_degen*.

by Figures 3.1 and 3.2 being identical).

The ideal parameter set is visually within 1σ of the best-fit in Figure 6.6, meaning that other parameter sets tested by *param_degen* with nearby masses and radii yield χ^2 that are up to 1σ away from the best-fit χ^2 . We will refer to this σ as the “*pd* σ ” (or “*param_degen* σ ”). In comparing the $\Delta\chi^2$ values of the ideal parameter set and the best-fit parameter set, the ideal is up to 2σ away from the best-fit (referred to as the “numerical σ ” or “num. σ ”). In theory, the *pd* σ and the numerical σ should be the same, but in this case and most of the following cases, we find that they are not. The difference between the *pd* σ and numerical σ suggests that *param_degen* is not finding all of the well-fitting parameter sets close to the best-fit, and improvements to this secondary

algorithm are discussed further in Section 6.11. We also notice that the best-fit has a lower χ^2 than the ideal. This suggests that the Poisson statistics, used to add low amounts of noise to the simulated data, shifted all the points by an overall trend or net amount, so that the ideal parameter set no longer yields the best-fitting light curve.

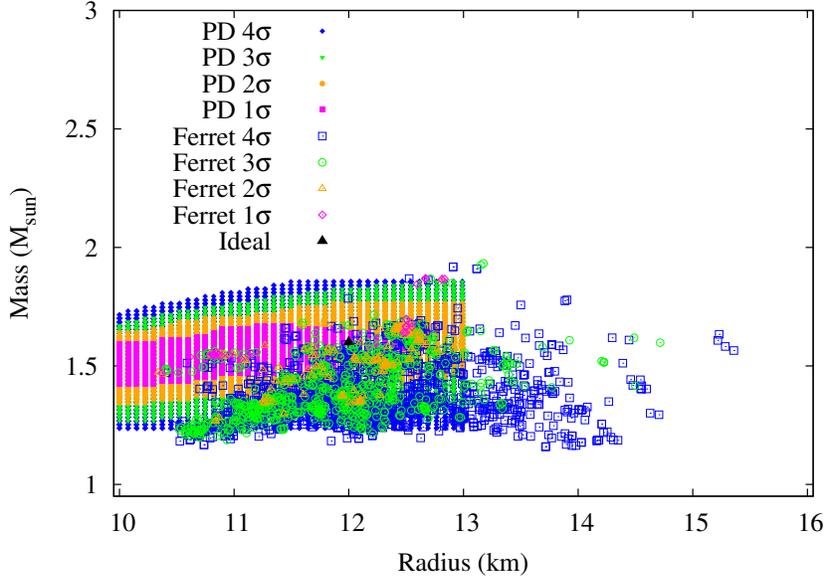


Figure 6.7: Case 1: A superposition of Figure 6.3 and Figure 6.6 for Fit 9. In the key, “PD” refers to the results of *param_degen*, and “Ferret” refers to the results of Ferret. We elect not to plot the 5σ regions here. We notice that the 1σ regions do not overlap. See discussion in text. Note that the y -axis extends up to $3M_{\odot}$ to fit the key for this plot.

The Ferret and *param_degen* contour regions do not overlap well in mass-radius space (Figure 6.7). This might be due to fixing the phase shift in *param_degen* to the Ferret best-fit phase shift value and not allowing it to vary.

Table 6.1 shows the best solution for all nine runs from Ferret for Case 1. The ideal solution is within 1σ for Fits 1, 2, 4, 6, 7, and 8.

In the subsequent cases, the ideal curves and fit curves in this section were created in the same way as with Case 1. For each “Fit” curve we use the parameter set of the lowest χ^2 for that case to make a Ferret χ^2 contour plot as

in Figure 6.3 and a parameter degeneracy χ^2 contour plot as in Figure 6.6. The ideal parameters and best-fit parameters for each of the ten cases are shown in Table 6.3. Results are discussed in detail in the next sections.

	χ^2	$M(M_\odot)$	% err _M	R (km)	% err _R	i (°)	θ (°)	$phase$	A	B	$pd \sigma$	Num. σ
Case 1	Ideal	58.88	1.60	12.00	-	10.00	60.00	0.0000	0.1317	0.0282	-	-
	Fit 1	57.26	1.64	12.67	2.5%	64.47	8.910	0.9957	0.1335	0.0275	-	1 σ
	Fit 2	56.67	1.65	12.88	3.1%	8.768	64.52	0.9959	0.1331	0.0274	-	1 σ
	Fit 3	56.11	1.63	11.81	1.9%	64.66	9.424	0.9964	0.1336	0.0276	-	2 σ
	Fit 4	58.46	1.66	13.73	3.8%	8.595	62.84	0.9972	0.1319	0.0277	-	1 σ
	Fit 5	55.55	1.55	11.24	3.1%	10.12	61.49	0.9978	0.1335	0.0275	-	2 σ
	Fit 6	57.33	1.59	11.21	0.6%	60.71	10.48	0.0010	0.1316	0.0284	-	1 σ
	Fit 7	57.42	1.64	10.72	2.5%	74.11	8.917	0.9914	0.1357	0.0263	-	1 σ
	Fit 8	58.18	1.44	11.46	10.0%	12.01	50.44	0.0028	0.1322	0.0282	-	1 σ
Fit 9	55.43	1.55	10.82	3.1%	60.97	10.59	0.9991	0.1331	0.0279	1 σ	2 σ	
Case 2	Ideal	49.88	1.60	14.00	-	40.00	60.00	0.0000	0.6236	0.1606	-	-
	Fit 1	47.02	1.60	13.82	0.0%	57.72	42.45	0.0010	0.6233	0.1631	-	2 σ
	Fit 2	47.19	1.57	13.60	1.9%	45.28	54.41	0.0016	0.6242	0.1622	-	2 σ
	Fit 3	45.39	1.60	13.81	0.0%	41.60	58.59	0.0008	0.6234	0.1617	4 σ	2 σ
Case 3	Ideal	72.14	1.40	13.00	-	10.00	10.00	0.0000	0.0210	0.0058	-	-
	Fit 1	70.67	1.79	13.43	27.9%	1.766	57.83	0.9844	0.0221	0.0055	-	1 σ
	Fit 2	70.56	1.72	11.73	22.9%	58.99	1.93	0.9836	0.0226	0.0055	1 σ	1 σ
Case 4	Ideal	66.18	1.50	13.00	-	60.00	20.00	0.0000	0.3001	0.0578	-	-
	Fit 1	60.43	1.44	12.03	4.0%	21.85	57.47	0.0006	0.3008	0.0571	-	2 σ
	Fit 2	60.03	1.34	11.56	10.7%	27.71	46.64	0.0045	0.2996	0.0587	-	2 σ
	Fit 3	58.32	1.36	11.22	9.3%	47.31	28.05	0.0048	0.2998	0.0589	5 σ	3 σ
Case 5	Ideal	44.95	1.50	13.00	-	60.00	20.00	0.0000	0.3001	0.0578	-	-
	Fit 1	43.78	1.42	11.81	5.3%	23.93	53.36	0.0032	0.2975	0.0582	4 σ	1 σ
	Fit 2	44.80	1.43	12.41	4.7%	22.90	53.90	0.0031	0.2974	0.0585	-	1 σ
	Fit 3	44.18	1.47	12.80	2.0%	20.27	59.42	0.0000	0.3007	0.0572	-	1 σ

<i>cont.</i>	χ^2	$M(M_\odot)$	% err _M	R (km)	% err _R	i (°)	θ (°)	<i>phase</i>	A	B	$pd \sigma$	Num. σ
Case 6 F ₂ , O	Ideal	55.51	-	13.00	-	60.00	20.00	0.0000	0.3001	0.0578	-	-
	Fit 1	45.16	5.3%	13.46	3.5%	55.10	21.14	0.0036	0.2997	0.0584	3 σ	3 σ
	Fit 2	45.61	7.3%	13.25	1.9%	53.34	21.85	0.0027	0.2979	0.0580	-	3 σ
	Fit 3	45.17	4.7%	12.85	1.2%	51.64	22.95	0.0038	0.2885	0.0584	-	3 σ
Case 7 F ₂ , G, O	Ideal	66.27	-	13.00	-	60.00	20.00	0.0000	0.3001	0.0578	-	-
	Fit 1	61.49	1.3%	12.62	2.9%	50.38	23.78	0.0042	0.2791	0.0590	N/A	2 σ
	Fit 2	62.49	0.7%	12.12	6.8%	40.39	29.51	0.0058	0.2565	0.0593	-	2 σ
	Fit 3	61.80	0.7%	12.74	2.0%	54.58	22.09	0.0026	0.2873	0.0588	-	2 σ
Case 8	Ideal	54.58	-	13.00	-	45.00	45.00	0.0000	0.4916	0.0975	-	-
	Fit 1	53.33	0.7%	13.11	0.9%	45.27	44.60	0.0003	0.4916	0.0980	-	1 σ
	Fit 2	53.13	1.3%	13.04	0.3%	49.23	41.27	0.0002	0.4913	0.0980	4 σ	1 σ
	Fit 3	53.47	1.3%	12.86	1.1%	47.92	42.92	0.0004	0.4913	0.0981	-	1 σ
Case 9	Ideal	58.45	-	12.00	-	70.00	15.00	0.0000	0.2481	0.0456	-	-
	Fit 1	58.80	3.8%	12.66	5.5%	73.66	13.66	0.9966	0.2497	0.0447	2 σ	1 σ
	Fit 2	59.78	7.5%	11.52	4.0%	62.54	17.07	0.0036	0.2476	0.0464	-	1 σ
	Fit 3	59.11	6.9%	11.37	5.3%	64.68	16.60	0.0019	0.2487	0.0456	-	1 σ
Case 10 F ₁ , N	Ideal	71.99	-	12.00	-	10.00	60.00	0.0000	0.1317	0.0282	-	-
	Fit 1	66.81	1.3%	10.32	14.0%	59.13	11.46	0.0025	0.1249	0.0293	-	2 σ
	Fit 2	66.78	9.4%	11.33	5.6%	68.70	9.32	0.9962	0.1263	0.0282	1 σ	2 σ

Table 6.3: Similar to Table 6.1. Models for each case have a spherical neutron star with isotropic emission and a 581 Hz spin frequency, unless otherwise specified. Key: F₁ = first set of fiducial parameters, F₂ = second set of fiducial parameters, O = oblate neutron star, G = graybody emission, N = high noise. Parameter degeneracy contour plots, for the “ $pd \sigma$ ” column, were only made for the lowest χ^2 fit of each case. “N/A” indicates that the ideal parameters do not have a $pd \sigma$ within the 5 σ or lower region of the best-fit. Figures to visually clarify the “ $pd \sigma$ ” results are in the following subsections. See text for further discussion.

6.2 Case 2: Strong Doppler Boost

The Case 2 light curve (Figure 6.8) is characterized by an asymmetric shape and a large pulse amplitude. The ideal model uses a spherical neutron star, isotropic emission, and a spin frequency of 800 Hz, with moderate inclination and high emission angles. The best-fit curve, using the parameters of Fit 3, is indistinguishable from the ideal curve in Figure 6.8.

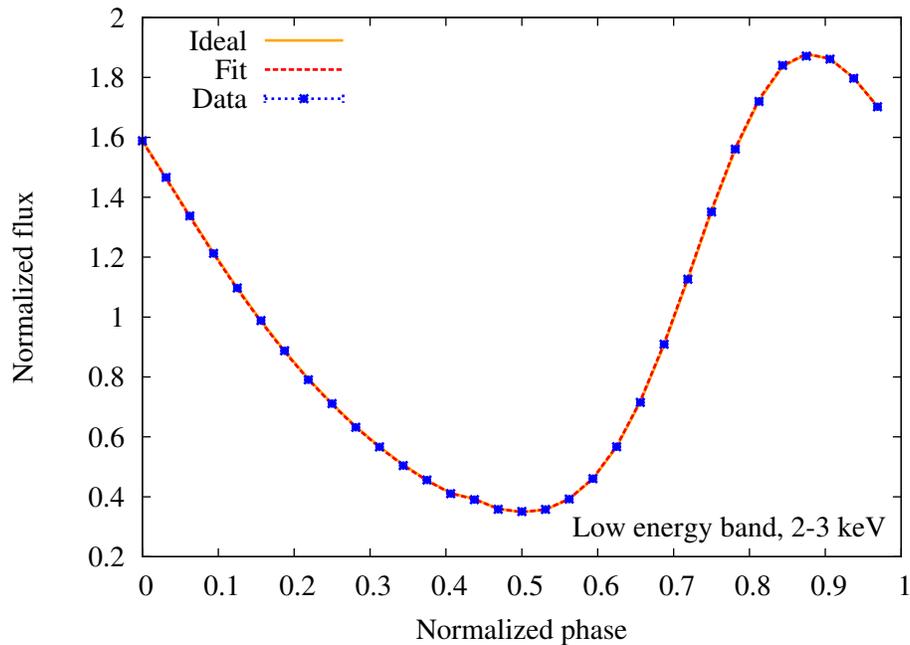


Figure 6.8: Case 2: ideal curve (orange solid), Ferret’s best-fit curve (red dotted), and simulated data (blue) for a spherical neutron star with isotropic emission, rotating at 800 Hz. Ideal parameters: $M = 1.6 M_{\odot}$, $R = 14$ km, $i = 40^{\circ}$, $\theta = 60^{\circ}$, phase shift = 0, $\chi^2 = 49.88$. Fit parameters: $M = 1.60 M_{\odot}$, $R = 13.81$ km, $i = 41.60^{\circ}$, $\theta = 58.59^{\circ}$, phase shift = 0.0008, $\chi^2 = 45.39$.

6.2.1 Ferret

The Ferret χ^2 contour plot (Figure 6.9 top panel) is very concentrated, showing that Ferret was easily able to find a large minima area, with four deeper minima within that. There is a small 3σ and 4σ island off to the right. The ideal

parameters fit nicely within the contour region. For Fit 3 the mass is 0% from the ideal and the radius is 1.4% from the ideal.

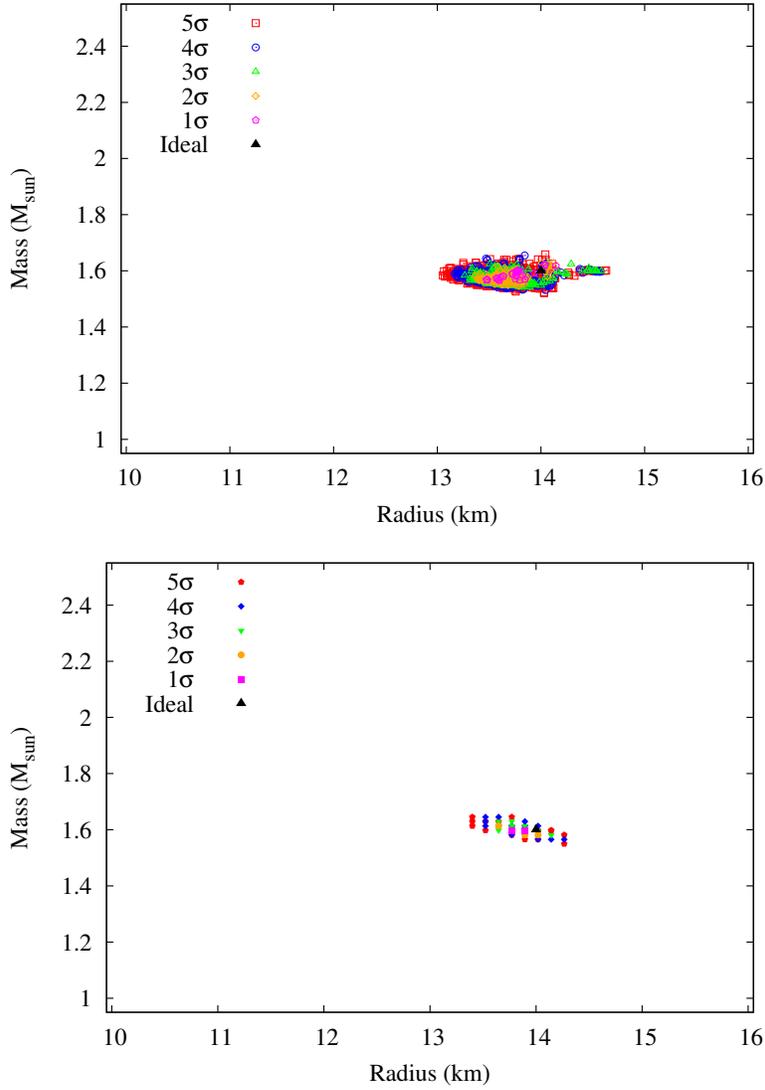


Figure 6.9: Case 2: Similar to Figures 6.3 and 6.6. The σ contour levels are calculated from Fit 3, $\chi^2 = 45.39$, which has $M = 1.60 M_\odot$ and $R = 13.81$ km. The χ^2 cut-offs are 47.68 for 1σ , 51.57 for 2σ , 57.21 for 3σ , 64.73 for 4σ , and 74.13 for 5σ . The ideal parameters are $M = 1.6 M_\odot$ and $R = 14$ km with $\chi^2 = 49.88$. The top panel shows the Ferret χ^2 contours. The bottom panel shows the *param_degen* χ^2 contours. The two χ^2 contour areas seem to be in approximately the same location in mass-radius space. As seen in the bottom panel, the ideal solution has a *pd* σ of 4σ from the best-fit.

6.2.2 Parameter Degeneracy

For Case 2 Fit 3 the σ contour area is a small oval (Figure 6.9 bottom panel). The different σ regions are not distinctly concentric or nested. The ideal solution has a pd σ of 4σ . Numerically, the ideal solution is within 2σ of the best-fit parameters.

6.3 Case 3: High Noise, Symmetric, Very Low Amplitude

The Case 3 light curve (Figure 6.10) is characterized by noisy data for a mostly symmetric shape with a very low pulse amplitude. The ideal model uses a spherical neutron star with isotropic emission, and low inclination and emission angles. The best-fit curve, using the parameters of Fit 2, is clearly different from the ideal curve in Figure 6.10. The parameter most notably off is the phase shift.

6.3.1 Ferret

The Ferret χ^2 contour plot (Figure 6.11 top panel) shows a very large minimum region. This is due to the large error bars and noisy statistics in the simulated data. The different contour regions appear to almost line up on the lower part of the graph, but are nested within each other on the upper part of the graph. The ideal parameters fit nicely within the contour region. For Fit 2 the mass is 22.9% from the ideal and the radius is 9.8% from the ideal.

6.3.2 Parameter Degeneracy

For Case 3 Fit 2 the σ contour area is massive, spanning the entirety of the allowed area (Figure 6.11 bottom panel). The diagonal cut-off line on the top left is likely from eliminating non-physical parameter sets. The ideal solution

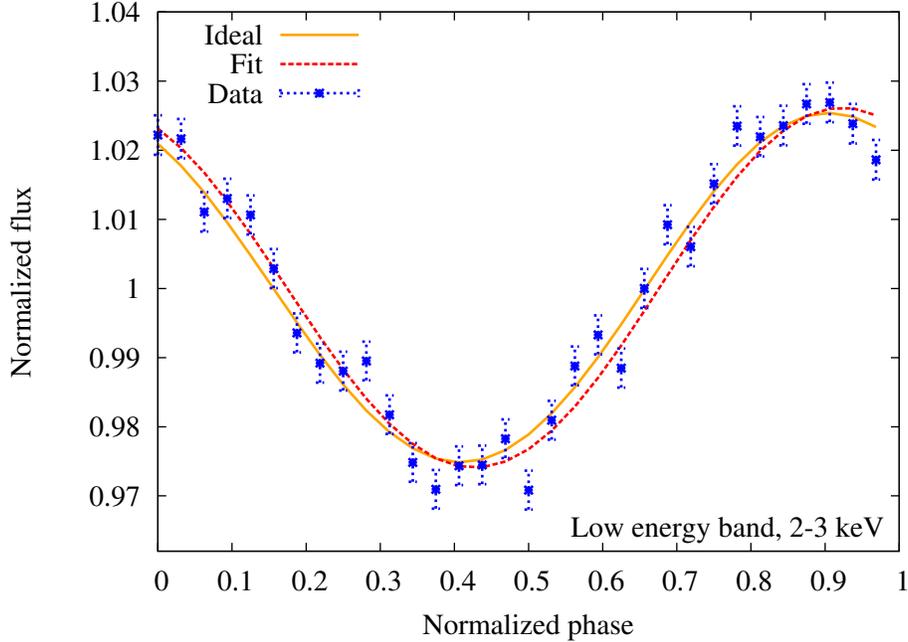


Figure 6.10: Case 3: ideal curve (orange solid), Ferret’s best-fit curve (red dotted), and simulated data (blue) for a spherical neutron star with isotropic emission, rotating at 581 Hz. Ideal parameters: $M = 1.4 M_{\odot}$, $R = 13$ km, $i = 10^{\circ}$, $\theta = 10^{\circ}$, phase shift = 0, $\chi^2 = 72.14$. Fit parameters: $M = 1.72 M_{\odot}$, $R = 11.73$ km, $i = 58.99^{\circ}$, $\theta = 1.93^{\circ}$, phase shift = 0.9836, $\chi^2 = 70.56$.

has a pd σ of 1σ , and numerically, the ideal solution is within 1σ of the best-fit parameters.

6.4 Case 4: Symmetric, Moderate Amplitude

The Case 4 light curve (Figure 6.12) is characterized by a symmetric shape and a moderate pulse amplitude. The ideal model uses a spherical neutron star with isotropic emission, and high inclination and low emission angles. The best-fit curve, using the parameters of Fit 3, is nearly identical to the ideal curve in Figure 6.12.

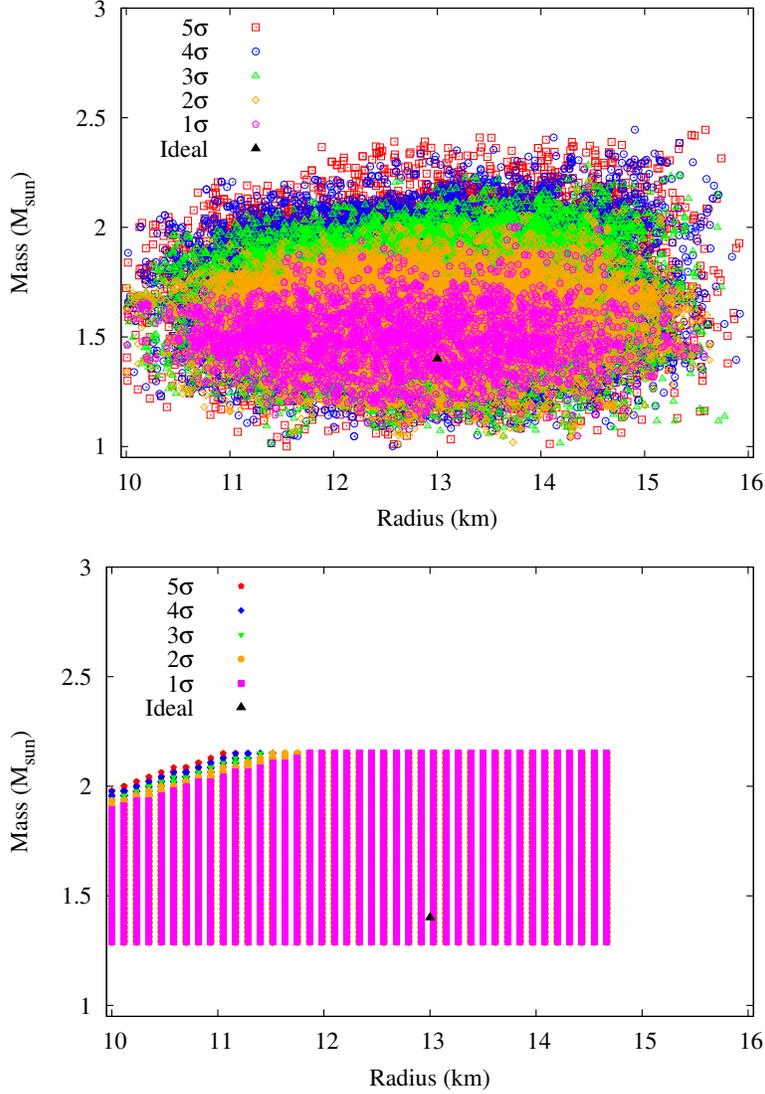


Figure 6.11: Case 3: Similar to Figure 6.9. The σ contour levels are calculated from Fit 2, $\chi^2 = 70.56$, which has $M = 1.72 M_{\odot}$ and $R = 11.73$ km. The χ^2 cut-offs are 72.85 for 1σ , 76.74 for 2σ , 82.39 for 3σ , 89.90 for 4σ , and 99.30 for 5σ . The ideal parameters are $M = 1.4 M_{\odot}$ and $R = 13$ km with $\chi^2 = 72.14$. As seen in the bottom panel, the ideal solution has a pd σ of 1σ from the best-fit. Note that the y -axis extends up to $3 M_{\odot}$ to fit the key for these plots. The abrupt cut-off in mass and radius in the bottom panel is due to setting the mass and radius range to be $\pm 25\%$ of the best-fit values for analysis with *param_degen*.

6.4.1 Ferret

The Ferret χ^2 contour plot (Figure 6.13 top panel) is more concentrated on the left, and less dense on the right. There appear to be three 1σ minima regions.

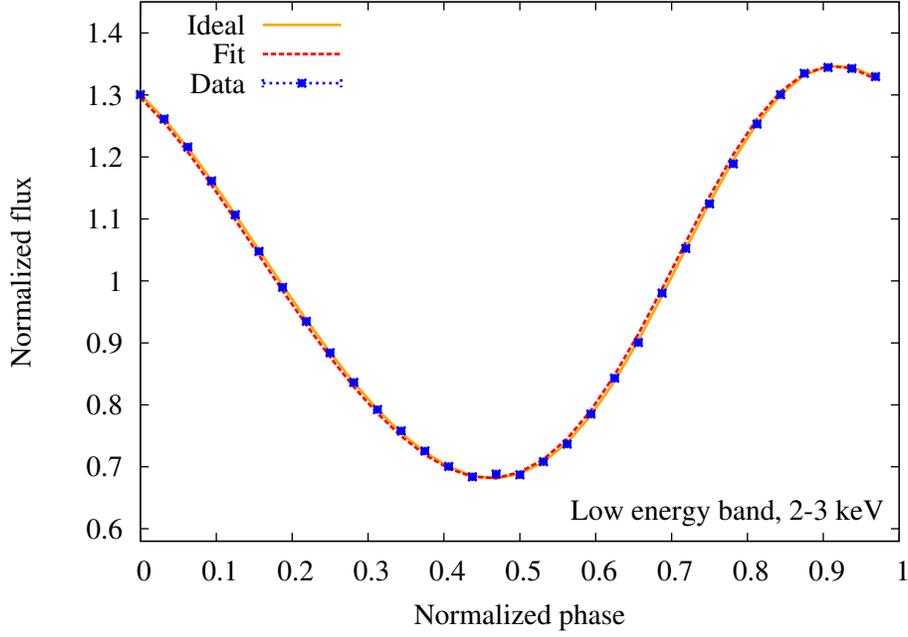


Figure 6.12: Case 4: ideal curve (orange solid), Ferret’s best-fit curve (red dotted), and simulated data (blue) for a spherical neutron star with isotropic emission, rotating at 581 Hz. Ideal parameters: $M = 1.5 M_{\odot}$, $R = 13$ km, $i = 60^{\circ}$, $\theta = 20^{\circ}$, phase shift = 0, $\chi^2 = 66.18$. Fit parameters: $M = 1.36 M_{\odot}$, $R = 11.22$ km, $i = 47.31^{\circ}$, $\theta = 28.05^{\circ}$, phase shift = 0.0048, $\chi^2 = 58.32$.

The ideal parameters are within the scattered region on the right. For Fit 3 the mass is 9.3% from the ideal and the radius is 12.7% from the ideal.

6.4.2 Parameter Degeneracy

For Case 4 Fit 3 the σ contour area is moderately large with defined but asymmetric σ regions (Figure 6.13 bottom panel). There are two 1σ islands within the primary 2σ region, and a small secondary 2σ on the top right. The ideal solution has a pd σ of 5σ , very close to the 4σ boundary. Numerically, the ideal solution is within 3σ of the best-fit parameters.

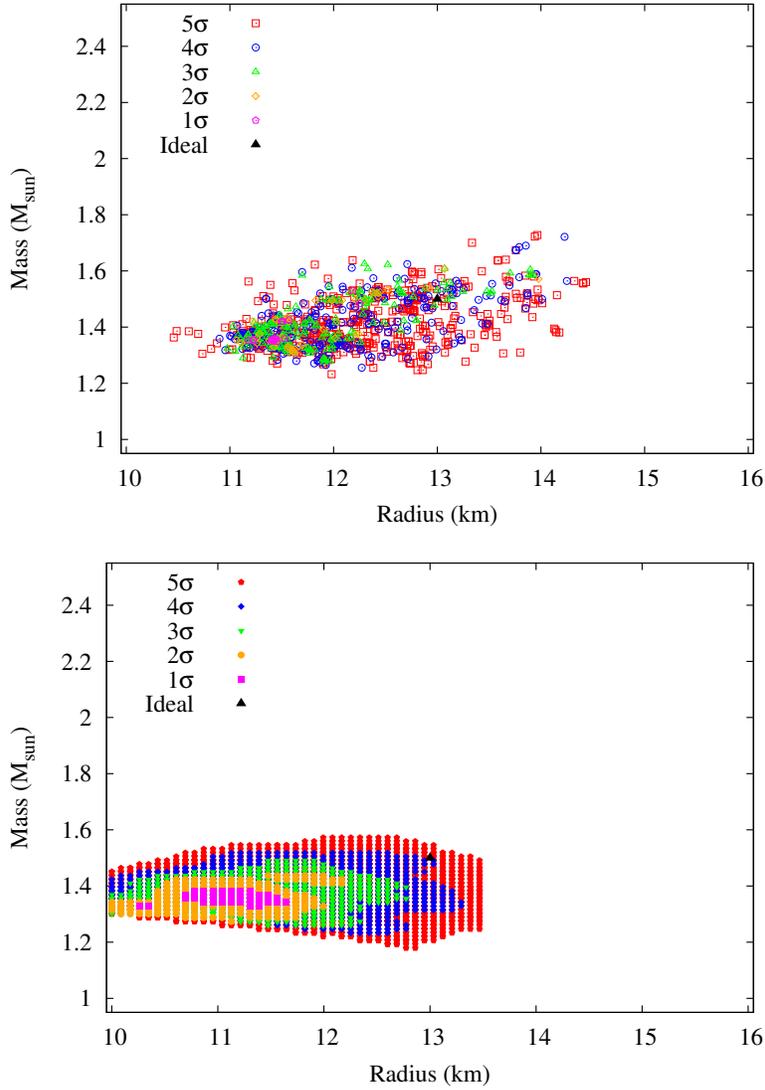


Figure 6.13: Case 4: Similar to Figure 6.9. The σ contour levels were calculated from Fit 3, $\chi^2 = 58.32$, which has $M = 1.36 M_{\odot}$ and $R = 11.22$ km. The χ^2 cut-offs are 60.62 for 1σ , 64.50 for 2σ , 70.15 for 3σ , 77.66 for 4σ , and 87.06 for 5σ . The ideal parameters are $M = 1.5 M_{\odot}$ and $R = 13$ km with $\chi^2 = 66.18$. As seen in the bottom panel, the ideal solution has a pd σ of 5σ from the best-fit.

6.5 Case 5: Graybody Model

The Case 5 light curve (Figure 6.14) is characterized by a fairly symmetric shape and a moderate pulse amplitude. The ideal model uses a spherical neutron star with graybody emission, and high inclination and low emission angles. It has

the same ideal parameters as Case 4. The best-fit curve, using the parameters of Fit 1, is nearly identical to the ideal curve in Figure 6.14.

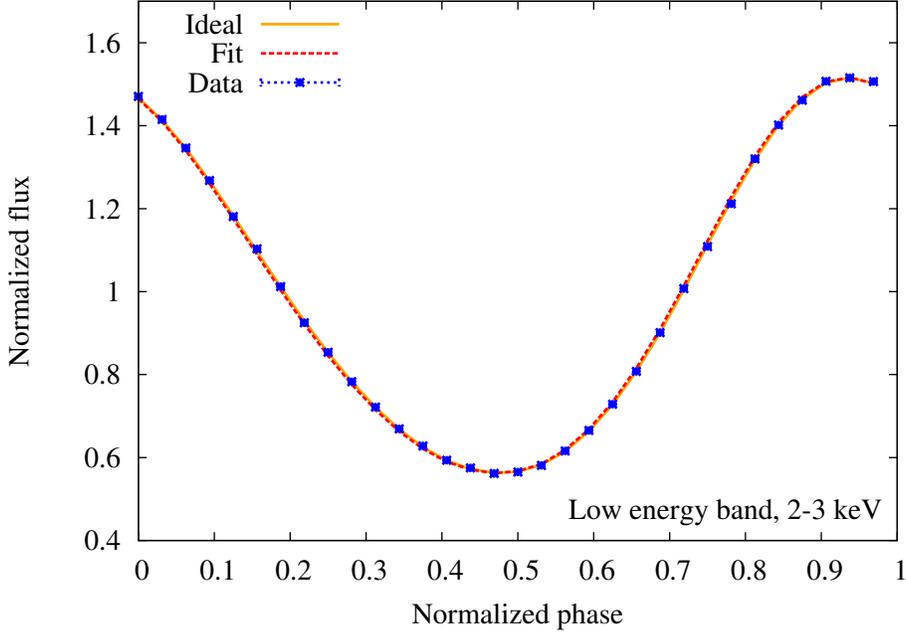


Figure 6.14: Case 5: ideal curve (orange solid), Ferret’s best-fit curve (red dotted), and simulated data (blue) for a spherical neutron star with gray-body emission, rotating at 581 Hz. Ideal parameters: $M = 1.5 M_{\odot}$, $R = 13$ km, $i = 60^{\circ}$, $\theta = 20^{\circ}$, phase shift = 0, $\chi^2 = 44.95$. Fit parameters: $M = 1.42 M_{\odot}$, $R = 11.81$ km, $i = 23.93^{\circ}$, $\theta = 53.36^{\circ}$, phase shift = 0.0032, $\chi^2 = 43.78$.

6.5.1 Ferret

The Ferret χ^2 contour plot (Figure 6.15 top panel) is very sparse. There are a few small areas that have more clustered individuals on the left, but the ideal parameters are in the very sparse scattered region on the right. For Fit 1 the mass is 5.3% from the ideal and the radius is 9.2% from the ideal.

6.5.2 Parameter Degeneracy

For Case 5 Fit 1 the σ contour regions are approximately concentric in a stretched out, eccentric ellipse shape, with one 1σ region in the middle (Fig-

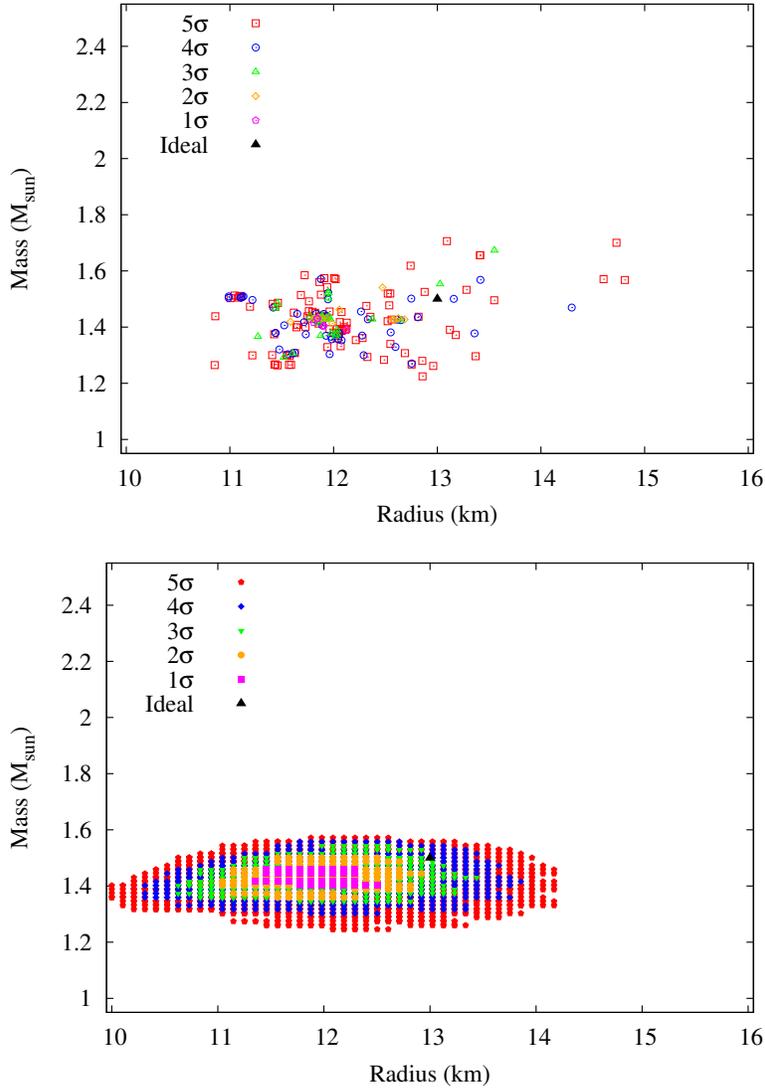


Figure 6.15: Case 5: Similar to Figure 6.9. The σ contour levels were calculated from Fit 1, $\chi^2 = 43.78$, which has $M = 1.42 M_{\odot}$ and $R = 11.81$ km. The χ^2 cut-offs are 46.07 for 1σ , 49.96 for 2σ , 55.61 for 3σ , 63.12 for 4σ , and 72.52 for 5σ . The ideal parameters are $M = 1.5 M_{\odot}$ and $R = 13$ km with $\chi^2 = 44.95$. As seen in the bottom panel, the ideal solution has a pd σ of 4σ from the best-fit.

ure 6.15 bottom panel). The ideal solution has a pd σ of 4σ , very close to the 3σ boundary. Numerically, the ideal solution is within 1σ of the best-fit parameters.

6.6 Case 6: Oblate Model

The Case 6 light curve (Figure 6.16) is characterized by a symmetric shape and a moderately low pulse amplitude. The ideal model uses an oblate neutron star with isotropic emission, and high inclination and low emission angles. It has the same parameters as Case 4. The best-fit curve, using the parameters of Fit 1, is nearly identical to the ideal curve in Figure 6.16, but the phase is slightly off.

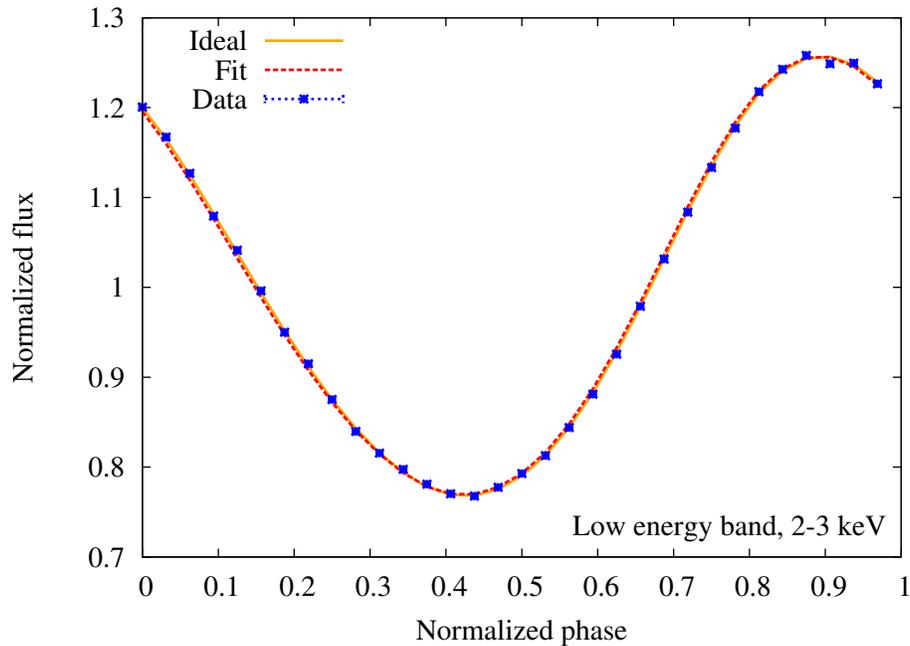


Figure 6.16: Case 6: ideal curve (orange solid), Ferret’s best-fit curve (red dotted), and simulated data (blue) for an oblate neutron star with isotropic emission, rotating at 581 Hz. Ideal parameters: $M = 1.5 M_{\odot}$, $R = 13$ km, $i = 60^{\circ}$, $\theta = 20^{\circ}$, phase shift = 0, $\chi^2 = 55.51$. Fit parameters: $M = 1.42 M_{\odot}$, $R = 13.46$ km, $i = 55.10^{\circ}$, $\theta = 21.14^{\circ}$, phase shift = 0.0036, $\chi^2 = 45.16$.

6.6.1 Ferret

The Ferret χ^2 contour plot (Figure 6.17 top panel) is quite concentrated in the middle. The contours are approximately nested, with two 1σ minima. The

ideal parameters fall within the clustered region, at $M = 1.5 M_{\odot}$ and $R = 13$ km. For Fit 1 the mass is 5.3% from the ideal and the radius is 3.5% from the ideal.

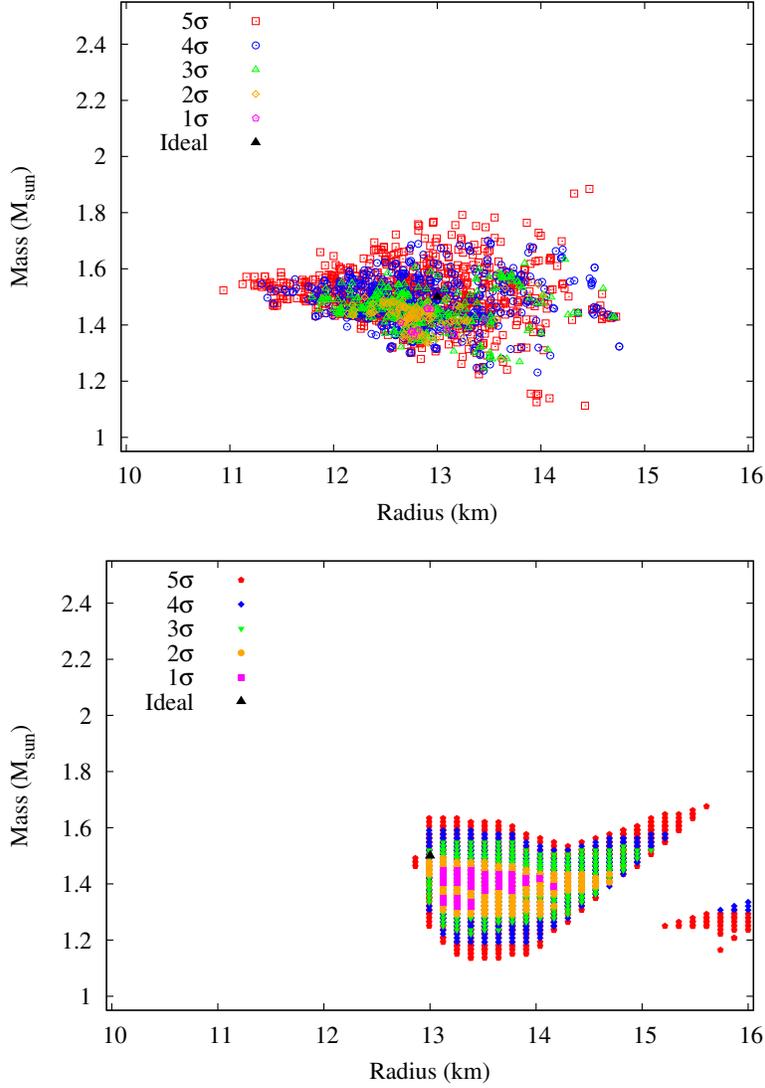


Figure 6.17: Case 6: Similar to Figure 6.9. The σ contour levels were calculated from Fit 1, $\chi^2 = 45.16$, which has $M = 1.42 M_{\odot}$ and $R = 13.46$ km. The χ^2 cut-offs are 47.45 for 1 σ , 51.34 for 2 σ , 56.99 for 3 σ , 64.50 for 4 σ , and 73.90 for 5 σ . The ideal solution has $M = 1.5 M_{\odot}$ and $R = 13$ km with $\chi^2 = 55.51$. As seen in the bottom panel, the ideal solution has a *pd* σ of 3 σ from the best-fit.

6.6.2 Parameter Degeneracy

For Case 6 Fit 1 there are two distinct sigma regions, and within the larger one, two distinct 1σ regions (Figure 6.17 bottom panel). The larger region is reminiscent of a very thick “v”. The ideal solution has a pd σ of 3σ , on the border of the 2σ region. Numerically, the ideal solution is within 3σ of the best-fit parameters.

6.7 Case 7: Oblate, Graybody Model

The Case 7 light curve (Figure 6.18) is characterized by a symmetric shape and a moderately low pulse amplitude. The ideal model uses an oblate neutron star with graybody emission, and high inclination and low emission angles. It has the same parameters as Case 4. In future light curve fitting with data we will also be using the oblate model with graybody emission, as this model is the most realistic of ours. The best-fit curve, using the parameters of Fit 1, is nearly indistinguishable from the ideal curve in Figure 6.18, but the phase shift is slightly off and the pulse amplitude is slightly larger.

6.7.1 Ferret

The Ferret χ^2 contour plot (Figure 6.19 top panel) is concentrated in a teardrop shape, with some outliers scattered out to higher masses and radii. The ideal parameters are near one of the 1σ minima in the clustered region. For Fit 1 the mass is 1.3% from the ideal and the radius is 2.9% from the ideal.

6.7.2 Parameter Degeneracy

For Case 7 Fit 1 there are two distinct σ regions (Figure 6.19 bottom panel). The first one is shaped a bit like an upside-down tear drop, with one 1σ region in the middle. The second σ region only has a small 5σ area. The ideal solution

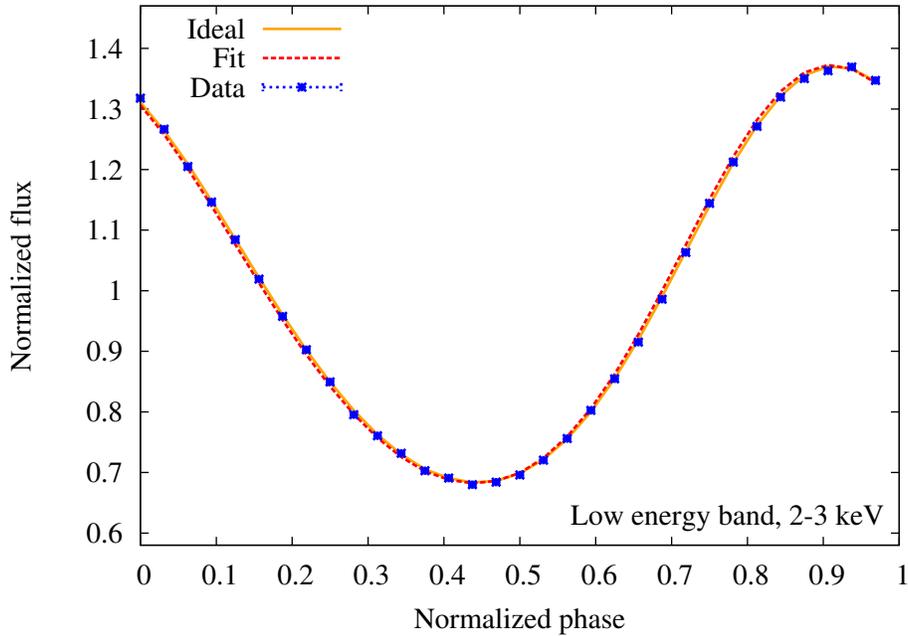


Figure 6.18: Case 7: ideal curve (orange solid), Ferret’s best-fit curve (red dotted), and simulated data (blue) for an oblate neutron star with graybody emission, rotating at 581 Hz. Ideal parameters: $M = 1.5 M_{\odot}$, $R = 13$ km, $i = 60^{\circ}$, $\theta = 20^{\circ}$, phase shift = 0, $\chi^2 = 66.27$. Fit parameters: $M = 1.48 M_{\odot}$, $R = 12.62$ km, $i = 50.38^{\circ}$, $\theta = 23.78^{\circ}$, phase shift = 0.0042, $\chi^2 = 61.49$.

has a pd σ greater than 5σ (denoted “N/A” in Table 6.3), but it is close to the 5σ edge. Numerically, the ideal solution is within 2σ of the best-fit parameters.

6.8 Case 8: Moderate Doppler Boost

The Case 8 light curve (Figure 6.20) is characterized by a somewhat asymmetric shape and a moderately high pulse amplitude. The ideal model uses a spherical neutron star with isotropic emission, and moderate inclination and emission angles of the same value. The best-fit curve, using the parameters of Fit 2, is identical to the ideal curve in Figure 6.20.

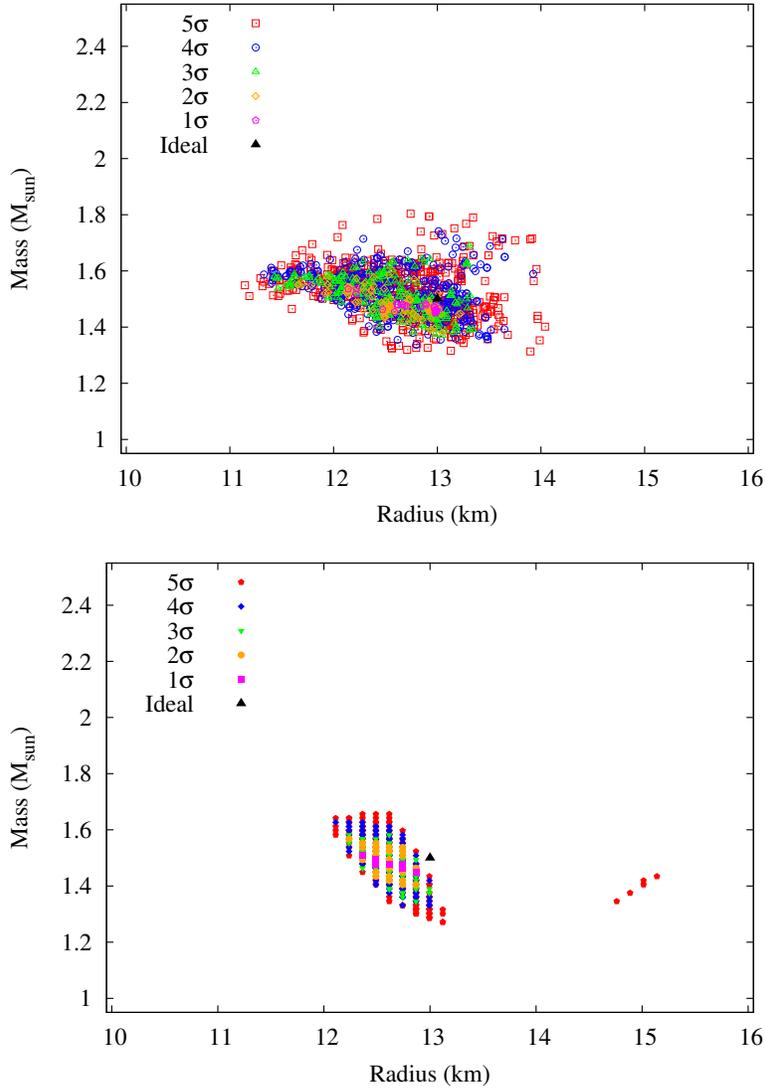


Figure 6.19: Case 7: Similar to Figure 6.9. The σ contour levels were calculated from Fit 1, $\chi^2 = 61.49$, which has $M = 1.48 M_{\odot}$ and $R = 12.62$ km. The χ^2 cut-offs are 63.78 for 1σ , 67.67 for 2σ , 73.32 for 3σ , 80.83 for 4σ , and 90.23 for 5σ . The ideal parameters are $M = 1.5 M_{\odot}$ and $R = 13$ km with $\chi^2 = 66.27$. As seen in the bottom panel, the ideal solution has a pd σ greater than 5σ from the best-fit.

6.8.1 Ferret

The Ferret χ^2 contour plot (Figure 6.21 top panel) is concentrated at moderate masses and radii, with some outlying points beyond $R \sim 14$ km and $M \sim 1.7 M_{\odot}$. The ideal parameters are comfortably in the 1σ – 3σ area in the clustered region.

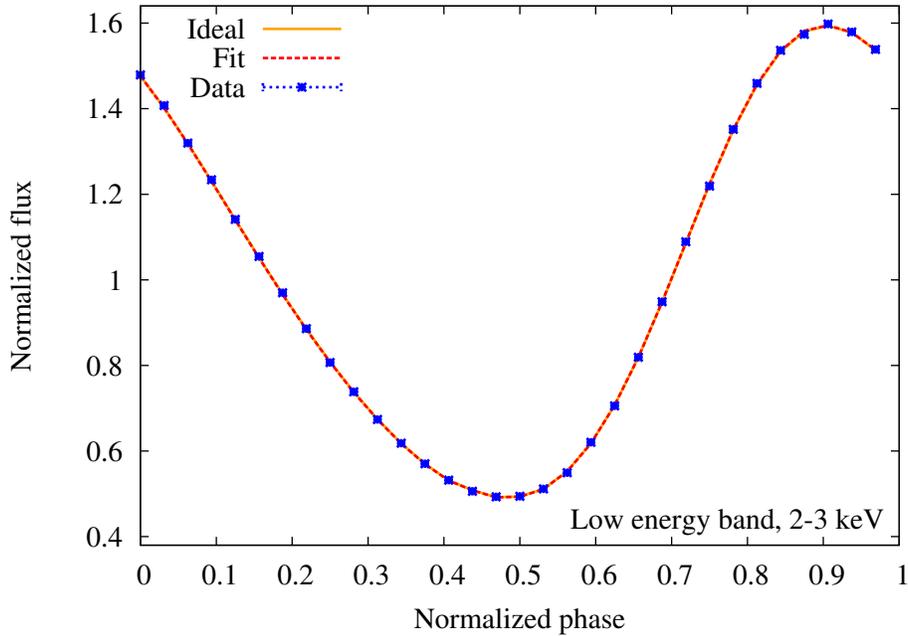


Figure 6.20: Case 8: ideal curve (orange solid), Ferret’s best-fit curve (red dotted), and simulated data (blue) for a spherical neutron star with isotropic emission, rotating at 581 Hz. Ideal parameters: $M = 1.5 M_{\odot}$, $R = 13$ km, $i = 45^{\circ}$, $\theta = 45^{\circ}$, phase shift = 0, $\chi^2 = 54.58$. Fit parameters: $M = 1.52 M_{\odot}$, $R = 13.04$ km, $i = 49.23^{\circ}$, $\theta = 41.27^{\circ}$, phase shift = 0.0002, $\chi^2 = 53.13$.

For Fit 2 the mass is 1.3% from the ideal and the radius is 0.3% from the ideal.

6.8.2 Parameter Degeneracy

For Case 8 Fit 2 the σ region looks a bit like a horizontal teardrop, with a little 5σ tail off the top (Figure 6.21 bottom panel). The σ regions are nested on the top, but approximately line up along the bottom edge. The ideal solution has a pd σ of 4σ . Numerically, the ideal solution is within 1σ of the best-fit parameters.

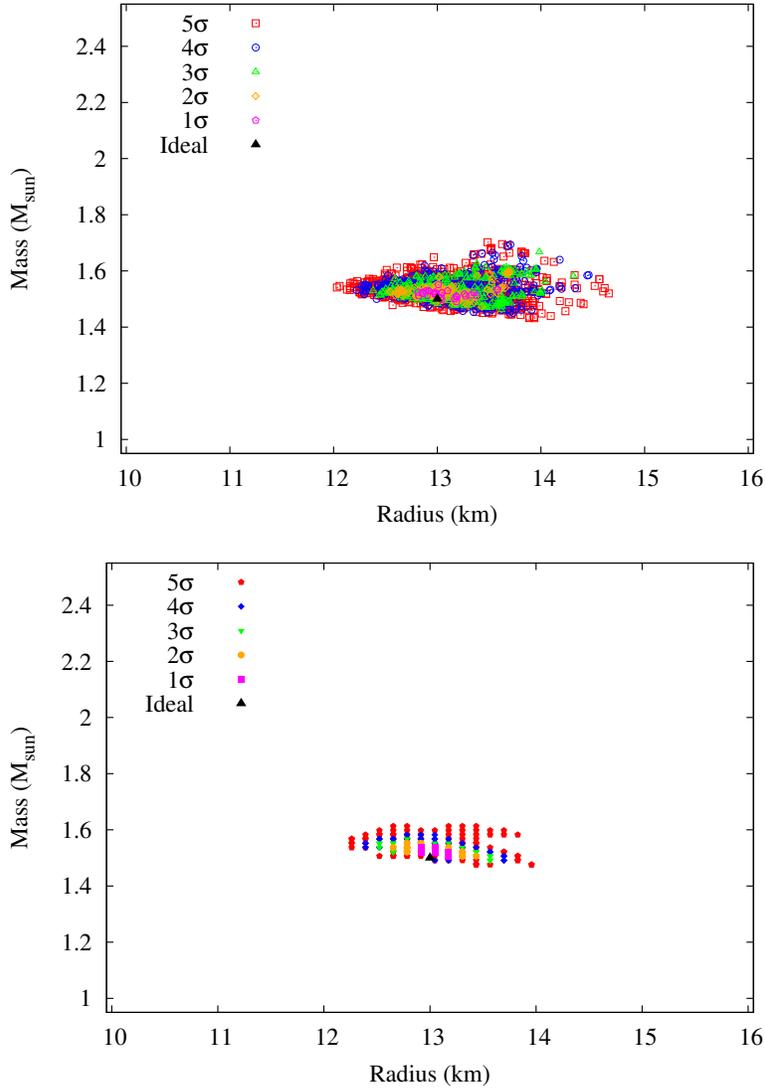


Figure 6.21: Case 8: Similar to Figure 6.9. The σ contour levels were calculated from Fit 2, $\chi^2 = 53.13$, which has $M = 1.52 M_{\odot}$ and $R = 13.04$ km. The χ^2 cut-offs are 55.43 for 1σ , 59.31 for 2σ , 64.96 for 3σ , 72.47 for 4σ , and 81.87 for 5σ . The ideal parameters are $M = 1.5 M_{\odot}$ and $R = 13$ km with $\chi^2 = 53.13$. As seen in the bottom panel, the ideal solution has a *pd* σ of 4σ from the best-fit.

6.9 Case 9: Symmetric, Moderately Low Amplitude

The Case 9 light curve (Figure 6.22) is characterized by a symmetric sinusoidal shape and a moderately low pulse amplitude. The ideal model uses a spherical

neutron star with isotropic emission, and very high inclination and low emission angles. The best-fit curve, using the parameters of Fit 1, is incredibly close to the ideal curve in Figure 6.22, but the phase shift is slightly off.

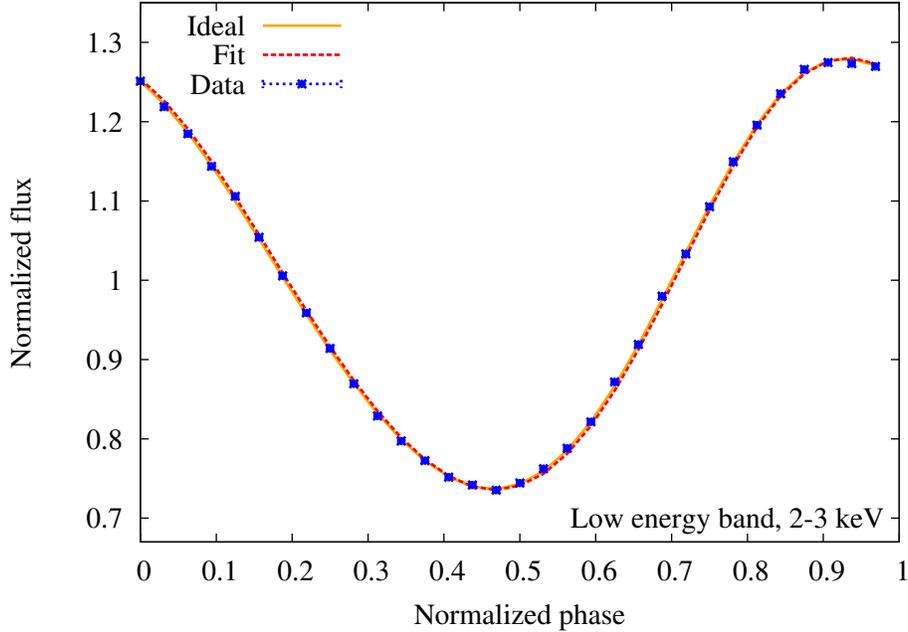


Figure 6.22: Case 9: ideal curve (orange solid), Ferret’s best-fit curve (red dotted), and simulated data (blue) for a spherical neutron star with isotropic emission, rotating at 581 Hz. Ideal parameters: $M = 1.6 M_{\odot}$, $R = 12$ km, $i = 70^{\circ}$, $\theta = 15^{\circ}$, phase shift = 0, $\chi^2 = 58.45$. Fit parameters: $M = 1.66 M_{\odot}$, $R = 12.66$ km, $i = 73.66^{\circ}$, $\theta = 13.66^{\circ}$, phase shift = 0.9966, $\chi^2 = 58.80$. This is the only case where the Ferret fit χ^2 is higher than the ideal χ^2 .

6.9.1 Ferret

The Ferret χ^2 contour plot (Figure 6.23 top panel) is very sparse and scattered, with a few small clustered areas. The ideal parameter set does not appear to be close to other parameter sets. Also note that this is the only case in which all three of Ferret’s best-fit χ^2 values were higher than the ideal χ^2 (see Table 6.3). For Fit 1 the mass is 3.8% from the ideal and the radius is 5.5% from the ideal.

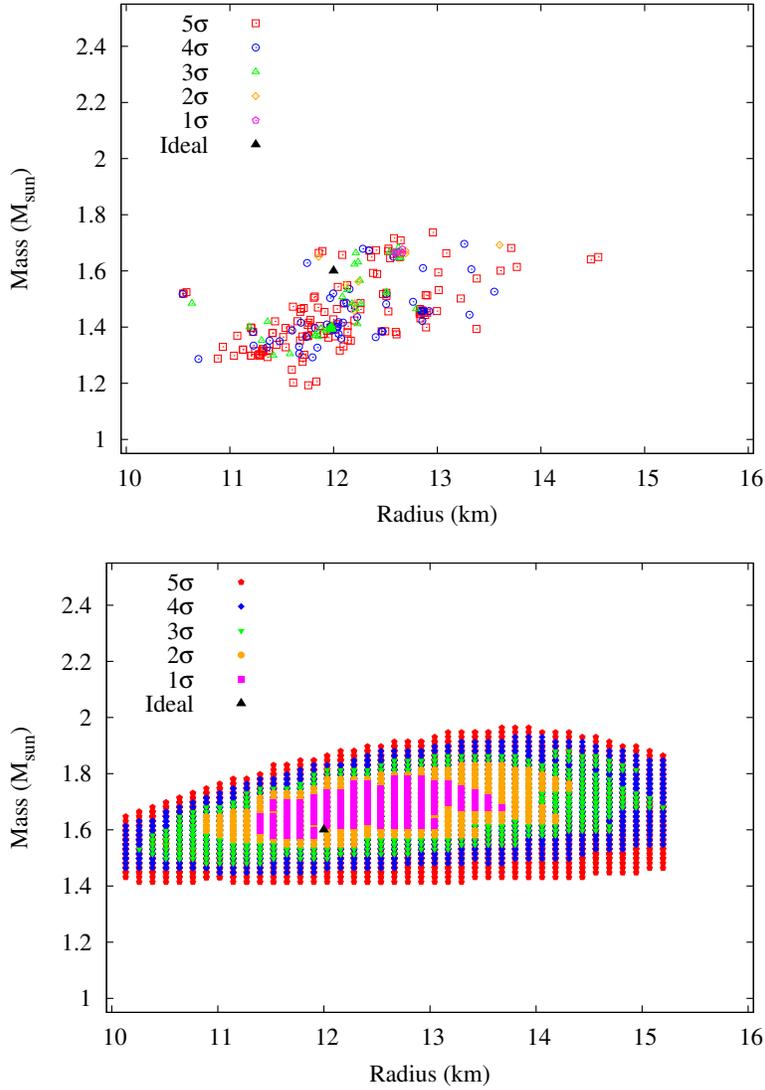


Figure 6.23: Case 9: Similar to Figure 6.9. The σ contour levels were calculated from Fit 1, $\chi^2 = 58.80$, which has $M = 1.66 M_\odot$ and $R = 12.66$ km. The χ^2 cut-offs are 61.10 for 1σ , 64.98 for 2σ , 70.63 for 3σ , 78.14 for 4σ , and 87.54 for 5σ . The ideal parameters are $M = 1.6 M_\odot$ and $R = 12$ km with $\chi^2 = 58.45$. This is the only case where the Ferret fit χ^2 is higher than the ideal χ^2 . As seen in the bottom panel, the ideal solution has a pd σ of 2σ from the best-fit.

6.9.2 Parameter Degeneracy

For Case 9 Fit 1 the large σ region spans all of the allowed radii, and has nested σ regions (Figure 6.23 bottom panel). The ideal solution has a pd σ of 2σ , close to the 1σ border. Numerically, the ideal solution is within 1σ of the best-fit

parameters.

6.10 Case 10: Fiducial, High Noise

The Case 10 light curve (Figure 6.24) is characterized by noisy data for a symmetric sinusoidal shape with a low pulse amplitude. The ideal model uses a spherical neutron star with isotropic emission, and low inclination and high emission angles. This has the same parameters as Case 1 but with noisy simulated data and large error bars. The best-fit curve, using the parameters of Fit 2, is distinguishable from the ideal curve in Figure 6.24. The pulse amplitude of the fit curve is a bit less than that of the ideal curve, but the phase shift appears to be correct.

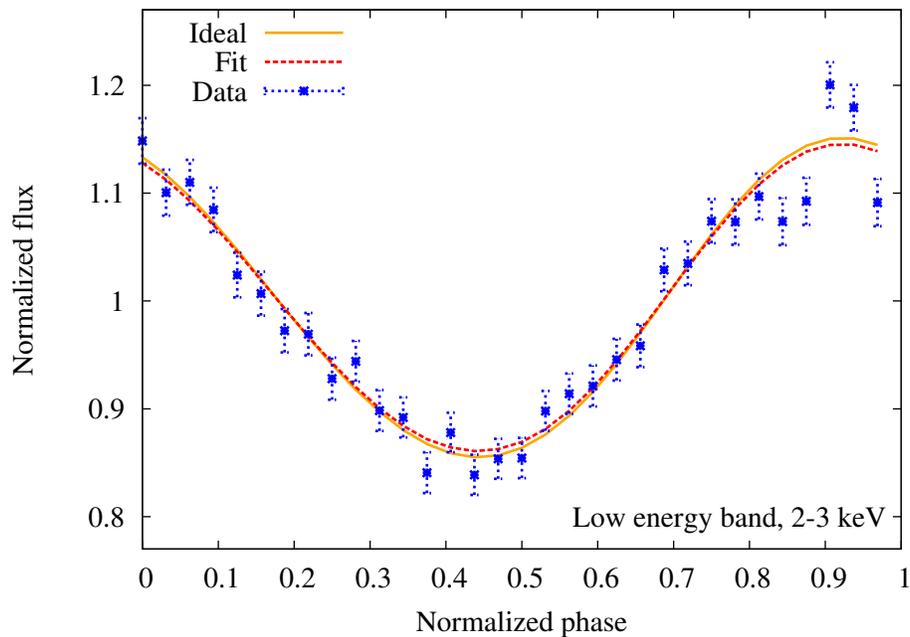


Figure 6.24: Case 10: ideal curve (orange solid), Ferret’s best-fit curve (red dotted), and simulated data (blue) for a spherical neutron star with isotropic emission, rotating at 581 Hz. Ideal parameters: $M = 1.6 M_{\odot}$, $R = 12$ km, $i = 10^{\circ}$, $\theta = 60^{\circ}$, phase shift = 0, $\chi^2 = 71.99$. Fit parameters: $M = 1.75 M_{\odot}$, $R = 11.33$ km, $i = 68.70^{\circ}$, $\theta = 9.32^{\circ}$, phase shift = 0.9962, $\chi^2 = 66.78$.

6.10.1 Ferret

The Ferret χ^2 contour plot (Figure 6.25 top panel) shows a very large concentric contour region. The ideal parameters are well within the 1σ area. For Fit 2 the mass is 9.4% from the ideal and the radius is 5.6% from the ideal.

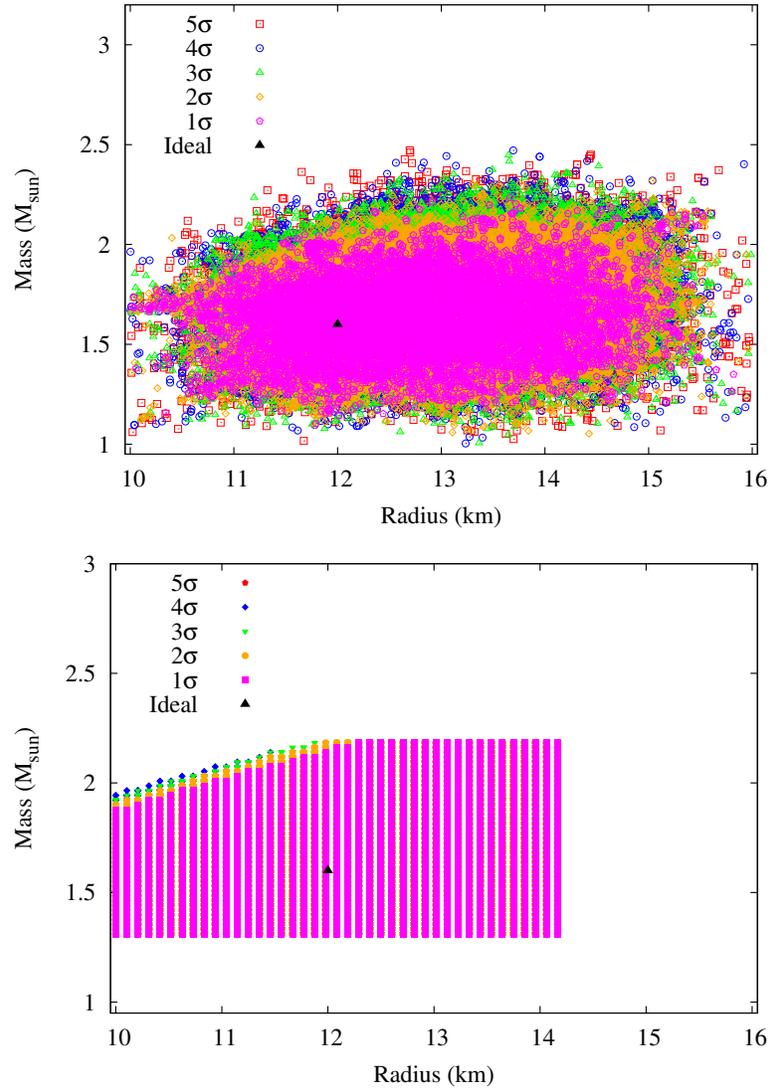


Figure 6.25: Case 10: Similar to Figure 6.9. The σ contour levels were calculated from Fit 2, $\chi^2 = 66.78$, which has $M = 1.75 M_{\odot}$ and $R = 11.33$ km. The χ^2 cut-offs are 69.07 for 1σ , 72.96 for 2σ , 78.61 for 3σ , 86.12 for 4σ , and 95.52 for 5σ . The ideal parameters are $M = 1.6 M_{\odot}$ and $R = 12$ km with $\chi^2 = 71.99$. As seen in the bottom panel, the ideal solution has a *pd* σ of 1σ from the best-fit. Note that the y -axis extends up to $3 M_{\odot}$ to fit the key for this plot.

6.10.2 Parameter Degeneracy

For Case 10 Fit 2 the σ contour area is massive (Figure 6.25 bottom panel). The diagonal cutoff line on top is from eliminating non-physical parameter sets, and the abrupt cut-offs in mass and radius come from only allowing a range of $\pm 25\%$ in mass and radius for analysis in *param_degen*. There is no visible 5σ region, likely due to the stacked points, where another parameter configuration with the same mass and radius yielded a lower χ^2 fit. The ideal solution has a *pd* σ of 1σ . Numerically, the ideal solution is within 2σ of the best-fit parameters.

6.11 Discussion and Summary

Overall, Ferret is able to fit the simulated data light curves well. The best accuracies in mass, at $< 1\%$ error, are for Case 2 Fit 1 and Case 2 Fit 3; the best accuracy in radius, at $< 1\%$ error, is for Case 8 Fit 2. Both Case 2 and Case 8 have low noise with small error bars, and moderate-to-strong Doppler effects. We propose that this detectable asymmetry in the light curve allows for such accurate fits, and for fitting observed light curve data we suggest the most accurate constraints will come from asymmetric curves. However the cause of this is not entirely clear. It was hypothesized that the more abrupt rise time seems to make the phase shifts line up better, since phase shift is the most sensitive parameter of the five. For example, a 5% deviation in phase shift will yield a much larger $\Delta\chi^2$ than a 5% deviation in mass, radius, inclination angle, or emission angle. As evident in Table 6.4, Case 2 does have the most sensitive phase shift deviation but Case 8 is not notably sensitive to phase shift, so this hypothesis does not have conclusive proof. Determining why asymmetry allows for better fits requires further investigation.

In addition, we see that the ideal parameter set is at most within the numerical 3σ range, as computed using $\Delta\chi^2$. Our goal in computing the *pd* σ is

Case	Δps for 1σ	Δps for 2σ
1	0.00146	0.00188
2	0.0000474	0.000111
3	0.00590	0.00862
4	0.000200	0.000400
5	0.000236	0.000392
6	0.00123	0.00144
7	0.000220	0.000416
8	0.000260	0.000385
9	0.000537	0.000799
10	0.00243	0.00527

Table 6.4: Sensitivity of χ^2 to deviations in phase shift for each case. Each phase shift discrepancy (Δps) was found by using the ideal parameters for each case, comparing with the simulated data, and then adjusting phase shift until the light curve gave a $\Delta\chi^2$ closest to $\Delta\chi^2 = 2.295$ for 1σ and $\Delta\chi^2 = 6.180$ for 2σ .

so that after adjustments to *param_degen*, the numerical σ is the same as *pd* σ for a simulated data set, and thus we will be able to give accurate mass and radius constraints out to different σ confidence contours. For these reasons, in the future when Ferret finds the best-fit parameters for a set of data, we feel it is safe to assume that the true mass and radius value is within the 3σ confidence region, as indicated by the numerical σ .

We also want to look at how degeneracies scale with noise. As shown in Table 6.3, the largest percent discrepancy of the best-fit mass and radius with the ideal mass and radius are 10.7% (Case 4 Fit 2) and 14.4% (Case 1 Fit 4), respectively, for data sets with low noise. Cases 1 and 4 have moderately low pulse amplitudes and appear very symmetric. It is likely that this symmetry allows for more parameter degeneracies, which yields a larger percent error. Of the high noise cases, the largest percent discrepancy in mass is 27.9% for Case 3 Fit 1, and the largest percent discrepancy in radius is 14.0% for Case 10 Fit 1. Cases 3 and 10 are the only two noisy curves. Case 3 has a very low pulse amplitude, which is a possible explanation as to why the percent discrepancy in mass was so high, since the Case 10 mass error was comparable to the low

noise cases. And yet, the Case 3 percent discrepancy in radius is comparable to Case 10 and the low noise cases, so there is not yet a definitive reason as to what allowed for such a high mass discrepancy for Case 3. The high noise cases require more analysis to draw strong conclusions.

We notice that the best-fit's χ^2 is lower than the ideal χ^2 fit with the simulated data, and moreover, the worse fits of the Ferret runs are those with a χ^2 closer to the ideal χ^2 . This likely comes from the Poisson statistics, used to noisify our simulated data, shifting the light curve by some net amount. However, it is potentially problematic: it implies that if we treat the ideal parameters as the true values of the neutron star, then the code presently is not able to determine the correct parameters.

Case 9 is the exception to the lower χ^2 problem: the best-fit χ^2 values are always higher than the ideal, even if only by a fraction. There is nothing that stands out about the Case 9 light curve in the low or high energy band that could explain this. It is presently unclear what the implications of this are. It is also possible that adjusting the Ferret setup parameters and allowing for more generations in a run will eliminate this problem.

Nevertheless, the fairly low percent error in mass and radius for the low noise cases is encouraging. We are collaborating more closely with J. D. Fiege in setting up the Ferret evolution parameters to yield better fits, and we are anticipating publishing the results of this further analysis.

One key observation that arises is that the *param_degen* σ contours and the numerical σ generally are not the same. For example, in Case 7 Fit 1, the ideal does not appear to be within 5σ from *param_degen*, but numerically it is within 2σ of the Ferret best-fit. This is probably due to the *param_degen* algorithm not being as thorough or nuanced as possible, and thus not accurately representing the σ boundaries in mass-radius space. We allow for some tolerance in A and B (equations (5.1.2) and (5.1.3)), as well as letting vary the mass, radius, inclination angle, and emission angle, but the one parameter not given any

leeway is the phase shift. While allowing it to vary will make the algorithm more computationally expensive, a small tolerance in phase shift with a moderately fine mesh may drastically improve the accuracy of the constraints.

The instances where the numerical σ is higher than the *param_degen* σ , in Case 1 and Case 10, are most confusing. These two cases have the same parameters and characteristics, with their only difference being noise levels in the simulated data, and they are the only two cases to exhibit this σ behaviour. It is possible that this effect may be due to an oversight in the *param_degen* algorithm, or may consistently arise from symmetric light curves with low pulse amplitudes, but the precise origin of this circumstance is unknown.

Furthermore, the *param_degen* algorithm currently does not have a way to account for graybody emission or oblateness in computing A and B . This makes the approximation much less precise. As seen in Figure 6.26, graybody and oblateness can distinctly change the light curve. For the given parameter set, graybody emission increases the pulse amplitude and causes the curve to peak a little later, whereas oblateness decreases the pulse amplitude, which makes it peak a little sooner in the normalized phase. When these two effects combine, as in the Case 7 curve, they appear to cancel, and so Case 4 and Case 7 are comparatively close. It also appears that, to first order, graybody emission and oblateness only affect pulse amplitude and not Doppler boosting, so we may only need to adjust A when improving the *param_degen* algorithm.

Allowing a tolerance in phase shift and accounting for graybody emission and oblateness in A and B probably will produce noticeable improvements in the mass and radius constraints that *param_degen* is able to make. We recommend directly computing the pulse amplitude — it is slightly more computationally expensive than the current approximation, but will give better results. Some tolerance in phase shift will likely give a larger σ area, which will not constrain the mass and radius fits of a neutron star as tightly, but we will be very confident in the constraints we will be able to provide.

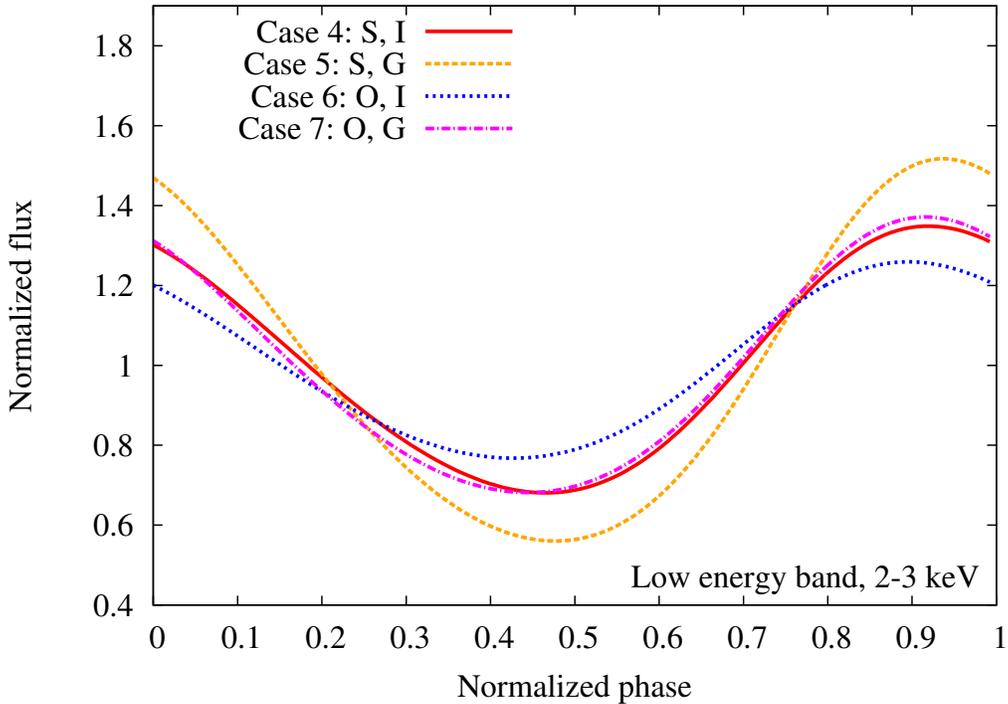


Figure 6.26: Ideal light curves for Cases 4, 5, 6, and 7, which have the same parameters: $M = 1.5 M_{\odot}$, $R = 13 \text{ km}$, $i = 60^{\circ}$, $\theta = 20^{\circ}$, phase shift = 0. As noted in the key, Case 4 (red solid) is a spherical neutron star with isotropic emission (S, I), Case 5 (orange dashed) is a spherical neutron star with graybody emission (S, G), Case 6 (blue dotted) is an oblate neutron star with isotropic emission (O, I), and Case 7 (magenta dash-dotted) is an oblate neutron star with graybody emission (O, G). We notice that for these parameters, graybody increases the pulse amplitude and oblateness decreases the pulse amplitude. These effects combine in the Case 7 light curve, so that the spherical isotropic model is comparatively close to the oblate graybody model.

For running with real data, we will use an oblate neutron star model with graybody emission such as in Case 7. While the best-fit does not currently include the ideal parameters within a pd σ of 5σ , the ideal is numerically within 2σ of all three fits so it is likely that with improvements to *param_degen* we will see more accurate mass and radius predictions and constraints.

In comparison with Lo et al. 2013 [39], we agree with their conclusion that systems with both a high inclination angle and high emission angle allow for tighter constraints on the predicted parameters, likely because these systems

have more Doppler boosting present in the light curves. Additionally, our findings do not disagree with their conclusion that low inclination angles and/or low emission angles yield light curves that are more difficult to fit with the correct parameters. The authors do not present the light curves for each of their models, so we are unable to directly compare light curve shapes with the constraints obtained.

Our results, as well as those published in the literature, substantiate the claim that parameter degeneracies are not due to the fitting method used, but are inherent in the problem of light curve fitting itself. Understanding the extent of these degeneracies is key to correctly interpreting the parameter constraints obtained from observed data.

Chapter 7

Conclusion

Determining the equation of state of supra-nuclear density cold matter is the holy grail of neutron star physics. The scientific community aims to do this by accurately and precisely measuring the masses and radii of a number of neutron stars. While other methods have yielded some advances in this (see Bhattacharyya 2010 [25] and Özel 2013 [90]), we presented a fitting procedure using our neutron star light curve model, which inputs the parameters of a neutron star and emitting hotspot region and outputs the observed light curve.

Our fitting procedure has two parts: the Ferret genetic algorithm, which fits the light curve data with our models using the “survival of the fittest” idea, and a parameter degeneracy constraint algorithm, which finds other parameter sets with the same light curve pulse amplitude and asymmetry as the best-fit from Ferret and plots them in mass-radius space so that we can determine the acceptable range of masses and radii within a given precision.

We implemented our fitting procedure by running through with ten cases of simulated data. Eight cases had low noise, comparable to the projected detection statistics for LOFT, and two cases had high noise, comparable to the detection statistics for RXTE. The fitting results show that we are able to accurately predict the mass of a neutron star to within at most 10.7% for low noise cases, and at most 27.9% for high noise cases; we are able to accurately

predict the radius of a neutron star to within at most 14.4% for low noise cases, and at most 14.0% for high noise cases. The χ^2 of the best-fit compared to the simulated data is within no more than 3σ from the χ^2 of the ideal parameter set compared with the simulated data. With improvements to the secondary parameter degeneracy algorithm, we are confident that we will be able to give appropriate mass and radius constraints for a data set based on the best-fit parameters from Ferret.

7.1 Future Research

Areas for future research include:

- Optimizing the *FerretSetup* parameters. We are currently in close collaboration with J. D. Fiege to ensure that the setup parameters are ideally suited to combat our parameter degeneracies, and we anticipate publishing the fitting results of this second round of analysis.
- Allowing some tolerance in the phase shift in the *param_degen* algorithm.
- Allowing the *Spot* code to model light curves from neutron stars with two non-antipodal hotspots.
- Determining if all oblate models give two separate σ contour regions in the parameter degeneracy plots as in the bottom panel in Figures 6.17 and 6.19.
- Determining whether the oblate neutron star light curve properties “cancel out” with the graybody emission light curve properties (as in Figure 6.26) for all parameter sets.
- Accommodating the simultaneous fitting of more energy bands, to incorporate some amount of spectral fitting.

- Looking at the effects of Poisson noise on the fits with multiple tests (each with different Poisson noise) of the same ideal light curve.
- Applying the Ferret genetic algorithm and the *param_degen* secondary algorithm to actual data from Type I X-ray bursts. Asymmetric light curves are preferable.
- Accounting for known parameter constraints and approximations for Type I X-ray burst sources. For example, if the inclination angle is approximately known for a system or if we have a mass approximation, this can be factored in to the Ferret genetic algorithm, and it may help to break some of the parameter degeneracies.
- Developing a Kerr metric approximation for an oblate neutron star, as suggested by Cadeau et al. 2007 [58].

With improvements like these, coupled with precise light curve measurements from a next-generation X-ray timing mission, we will be able to place constraints on the masses and radii of neutron stars to help deduce the equation of state for cold ultra-dense matter.

Appendix A

This appendix shows a hotspot temperature mesh to create a “tall” hotspot shape, as described in Section 3.4. Each value shown in the table is the temperature value for that cell over the hotspot. A kT temperature of 0 keV treats that cell as a non-emitting background, and a non-zero temperature treats that cell as an emitting region.

0	0	0	0	0	0	0	0	2	0	0	0	0	0	0
0	0	0	0	0	0	0	0	2	0	0	0	0	0	0
0	0	0	0	0	0	0	0	2	0	0	0	0	0	0
0	0	0	0	0	0	0	0	2	0	0	0	0	0	0
0	0	0	0	0	0	0	0	2	0	0	0	0	0	0
0	0	0	0	0	0	0	0	2	0	0	0	0	0	0
0	0	0	0	0	0	0	0	2	0	0	0	0	0	0
0	0	0	0	0	0	0	0	2	0	0	0	0	0	0
0	0	0	0	0	0	0	0	2	0	0	0	0	0	0
0	0	0	0	0	0	0	0	2	0	0	0	0	0	0
0	0	0	0	0	0	0	0	2	0	0	0	0	0	0
0	0	0	0	0	0	0	0	2	0	0	0	0	0	0
0	0	0	0	0	0	0	0	2	0	0	0	0	0	0
0	0	0	0	0	0	0	0	2	0	0	0	0	0	0
0	0	0	0	0	0	0	0	2	0	0	0	0	0	0
0	0	0	0	0	0	0	0	2	0	0	0	0	0	0
0	0	0	0	0	0	0	0	2	0	0	0	0	0	0
0	0	0	0	0	0	0	0	2	0	0	0	0	0	0

Table A.1: A 15×15 temperature mesh for a hotspot with $\rho = 15^\circ$, so that each cell shown here is 2° in diameter. See Figure 3.3 for approximate cell shapes.

Bibliography

- [1] N. K. Glendenning. *Compact Stars: Nuclear Physics, Particle Physics, and General Relativity*. Astronomy and Astrophysics Library. Springer-Verlag, 2nd edition, 2000.
- [2] A. Hewish, S. J. Bell, J. D. H. Pilkington, P. F. Scott, and R. A. Collins. Observation of a Rapidly Pulsating Radio Source. *Nat.*, 217:709–713, February 1968. doi: 10.1038/217709a0.
- [3] N. E. White, F. Nagase, and A. N. Parmar. The properties of x-ray binaries. In W. G. H. Lewin, J. van Paradijs, and E. P. J. van den Heuvel, editors, *X-Ray Binaries*, Cambridge Astrophysics Series. Cambridge University Press, 1995.
- [4] S. L. Shapiro and S. A. Teukolsky. *Black Holes, White Dwarfs, and Neutron Stars: The Physics of Compact Objects*. John Wiley & Sons, 1983.
- [5] M. Camenzind. *Compact Objects in Astrophysics: White Dwarfs, Neutron Stars, and Black Holes*. Astronomy and Astrophysics Library. Springer-Verlag, 2007.
- [6] J. R. Oppenheimer and G. M. Volkoff. On Massive Neutron Cores. *Phys. Rev.*, 55:374–381, February 1939. doi: 10.1103/PhysRev.55.374.
- [7] B. W. Carroll and D. A. Ostlie. *An Introduction to Modern Astrophysics*. Pearson Addison Wesley, 2nd edition, 2007.

- [8] K. G. Elshamouty, C. O. Heinke, G. R. Sivakoff, W. C. G. Ho, P. S. Shternin, D. G. Yakovlev, D. J. Patnaude, and L. David. Measuring the Cooling of the Neutron Star in Cassiopeia A with all Chandra X-ray Observatory Detectors. *Submitted. ArXiv e-prints: 1306.3387*, June 2013.
- [9] W. C. G. Ho, N. Andersson, C. M. Espinoza, K. Glampedakis, B. Haskell, and C. O. Heinke. The hottest superfluid and superconductor in the Universe: Discovery and nuclear physics implications. *Submitted. ArXiv e-prints: 1303.3282*, March 2013.
- [10] G. B. Cook, S. L. Shapiro, and S. A. Teukolsky. Rapidly rotating neutron stars in general relativity: Realistic equations of state. *ApJ*, 424:823–845, April 1994. doi: 10.1086/173934.
- [11] N. Ivanova. Astro 465 class lectures [*unpublished*], Winter 2012.
- [12] P. Ghosh and F. K. Lamb. Accretion by rotating magnetic neutron stars. ii - radial and vertical structure of the transition zone in disk accretion. *ApJ*, 232:259–276, August 1979.
- [13] R. Wijnands and M. van der Klis. A millisecond pulsar in an X-ray binary system. *Nat.*, 394:344–346, July 1998. doi: 10.1038/28557.
- [14] A. Papitto, T. Di Salvo, L. Burderi, T. M. Belloni, L. Stella, E. Bozzo, A. D’Aì, C. Ferrigno, R. Iaria, S. Motta, A. Riggio, and A. Tramacere. The pulse profile and spin evolution of the accreting pulsar in Terzan 5, IGR J17480-2446, during its 2010 outburst. *MNRAS*, 423:1178–1193, June 2012. doi: 10.1111/j.1365-2966.2012.20945.x.
- [15] G. W. Clark. X-ray binaries in globular clusters. *ApJL*, 199:L143–L145, August 1975. doi: 10.1086/181869.
- [16] A. Papitto, C. Ferrigno, E. Bozzo, N. Rea, L. Pavan, S. Campana, P. Romano, L. Burderi, T. Di Salvo, A. Riggio, D. F. Torres, M. Falanga,

- J. W. T. Hessels, M. Burgay, J. M. Sarkissian, M. H. Wieringa, M. D. Filipović, and G. F. Wong. Swinging between rotation and accretion power in a millisecond binary pulsar. *Submitted. ArXiv e-prints: 1305.3884*, May 2013.
- [17] W. H. G. Lewin, J. van Paradijs, and R. E. Taam. X-Ray Bursts. *Space Sci. Rev.*, 62:223–389, September 1993. doi: 10.1007/BF00196124.
- [18] A. L. Watts. Thermonuclear Burst Oscillations. *ARA&A*, 50:609–640, September 2012. doi: 10.1146/annurev-astro-040312-132617.
- [19] A. Cumming. Thermonuclear X-ray bursts: theory vs. observations. *NuPhS*, 132:435–445, June 2004. doi: 10.1016/j.nuclphysbps.2004.04.078.
- [20] R. Artigue, D. Barret, F. K. Lamb, K. H. Lo, and M. C. Miller. Testing the rotating hotspot model using X-ray burst oscillations from 4U 1636-536. *MNRAS*, 433:L64–L68, June 2013. doi: 10.1093/mnrasl/slt059.
- [21] A. Cumming, E. Zweibel, and L. Bildsten. Magnetic Screening in Accreting Neutron Stars. *ApJ*, 557:958–966, August 2001. doi: 10.1086/321658.
- [22] M. P. Muno, F. Özel, and D. Chakrabarty. The Amplitude Evolution and Harmonic Content of Millisecond Oscillations in Thermonuclear X-Ray Bursts. *ApJ*, 581:550–561, December 2002. doi: 10.1086/344152.
- [23] T. E. Strohmayer, W. Zhang, and J. H. Swank. 363 HZ Oscillations during the Rising Phase of Bursts from 4U 1728-34: Evidence for Rotational Modulation. *ApJL*, 487:L77, September 1997. doi: 10.1086/310880.
- [24] K. R. Pechenick, C. Ftaclas, and J. M. Cohen. Hot spots on neutron stars - The near-field gravitational lens. *ApJ*, 274:846–857, November 1983. doi: 10.1086/161498.

- [25] S. Bhattacharyya. Measurement of neutron star parameters: A review of methods for low-mass X-ray binaries. *Adv. Space Res.*, 45:949–978, April 2010. doi: 10.1016/j.asr.2010.01.010.
- [26] D. A. Leahy, S. M. Morsink, and Y. Chou. Constraints on the Mass and Radius of the Neutron Star XTE J1807-294. *ApJ*, 742:17, November 2011. doi: 10.1088/0004-637X/742/1/17.
- [27] S. M. Morsink and D. A. Leahy. Multi-epoch Analysis of Pulse Shapes from the Neutron Star SAX J1808.4-3658. *ApJ*, 726:56, January 2011. doi: 10.1088/0004-637X/726/1/56.
- [28] A. K. Harding and A. G. Muslimov. Pulsar X-Ray and Gamma-Ray Pulse Profiles: Constraint on Obliquity and Observer Angles. *ApJ*, 500:862, June 1998. doi: 10.1086/305763.
- [29] S. DeDeo, D. Psaltis, and R. Narayan. General Relativistic Constraints on Emission Models of Anomalous X-Ray Pulsars. *ApJ*, 559:346–352, September 2001. doi: 10.1086/322283.
- [30] R. Perna and E. V. Gotthelf. Constraints on the Emission and Viewing Geometry of the Transient Anomalous X-Ray Pulsar XTE J1810-197. *ApJ*, 681:522–529, July 2008. doi: 10.1086/588211.
- [31] U. Kraus. Hollow Accretion Columns on Neutron Stars and the Effects of Gravitational Light Bending. *ApJ*, 563:289–300, December 2001. doi: 10.1086/323791.
- [32] D. A. Leahy. Mass-Radius Constraints from a Pulse Shape Model for Hercules X-1. *ApJ*, 613:517–521, September 2004. doi: 10.1086/422905.
- [33] M. Sasaki, D. Klochkov, U. Kraus, I. Caballero, and A. Santangelo. Analyzing X-ray pulsar profiles: geometry and beam pattern of EXO 2030+375. *A&A*, 517:A8, July 2010. doi: 10.1051/0004-6361/200913863.

- [34] M. Annala and J. Poutanen. Constraining compactness and magnetic field geometry of X-ray pulsars from the statistics of their pulse profiles. *A&A*, 520:A76, September 2010. doi: 10.1051/0004-6361/200912773.
- [35] I. Caballero, U. Kraus, A. Santangelo, M. Sasaki, and P. Kretschmar. Analyzing X-ray pulsar profiles: geometry and beam pattern of A 0535+26. *A&A*, 526:A131, February 2011. doi: 10.1051/0004-6361/201014728.
- [36] D. Psaltis and T. Johannsen. A Ray-tracing Algorithm for Spinning Compact Object Spacetimes with Arbitrary Quadrupole Moments. I. Quasi-Kerr Black Holes. *ApJ*, 745:1, January 2012. doi: 10.1088/0004-637X/745/1/1.
- [37] C. K. Chan, D. Psaltis, and F. Ozel. GRay: a Massively Parallel GPU-Based Code for Ray Tracing in Relativistic Spacetimes. *Submitted. ArXiv e-prints: 1303.5057*, March 2013.
- [38] N. R. Nath, T. E. Strohmayer, and J. H. Swank. Bounds on Compactness for Low-Mass X-Ray Binary Neutron Stars from X-Ray Burst Oscillations. *ApJ*, 564:353–360, January 2002. doi: 10.1086/324132.
- [39] K. H. Lo, M. C. Miller, S. Bhattacharyya, and F. K. Lamb. Determining neutron star masses and radii using energy-resolved waveforms of X-ray burst oscillations. *Submitted. ArXiv e-prints: 1304.2330*, April 2013.
- [40] N. Weinberg, M. C. Miller, and D. Q. Lamb. Oscillation Waveforms and Amplitudes from Hot Spots on Neutron Stars. *ApJ*, 546:1098–1106, January 2001. doi: 10.1086/318279.
- [41] S. Bogdanov, G. B. Rybicki, and J. E. Grindlay. Constraints on Neutron Star Properties from X-Ray Observations of Millisecond Pulsars. *ApJ*, 670:668–676, November 2007. doi: 10.1086/520793.

- [42] S. Bogdanov. The Nearest Millisecond Pulsar Revisited with XMM-Newton: Improved Mass-radius Constraints for PSR J0437-4715. *ApJ*, 762:96, January 2013. doi: 10.1088/0004-637X/762/2/96.
- [43] P. Meszaros and H. Riffert. Gravitational light bending near neutron stars. II - Accreting pulsar spectra as a function of phase. *ApJ*, 327:712–722, April 1988. doi: 10.1086/166227.
- [44] M. P. Muno, F. Özel, and D. Chakrabarty. The Energy Dependence of Millisecond Oscillations in Thermonuclear X-Ray Bursts. *ApJ*, 595:1066–1076, October 2003. doi: 10.1086/377447.
- [45] S. Bhattacharyya, T. E. Strohmayer, M. C. Miller, and C. B. Markwardt. Constraints on Neutron Star Parameters from Burst Oscillation Light Curves of the Accreting Millisecond Pulsar XTE J1814-338. *ApJ*, 619:483–491, January 2005. doi: 10.1086/426383.
- [46] W. C. G. Ho. Constraining the geometry of the neutron star RX J1856.5-3754. *MNRAS*, 380:71–77, September 2007. doi: 10.1111/j.1365-2966.2007.12043.x.
- [47] E. V. Gotthelf, R. Perna, and J. P. Halpern. Modeling the Surface X-ray Emission and Viewing Geometry of PSR J0821-4300 in Puppis A. *ApJ*, 724:1316–1324, December 2010. doi: 10.1088/0004-637X/724/2/1316.
- [48] T. E. Strohmayer. Light curves of rotating, oscillating neutron stars. *ApJ*, 388:138–147, March 1992. doi: 10.1086/171136.
- [49] U. Lee and T. E. Strohmayer. Light curves of oscillating neutron stars. *MNRAS*, 361:659–672, August 2005. doi: 10.1111/j.1365-2966.2005.09198.x.
- [50] K. Numata and U. Lee. Light curves from rapidly rotating neutron

- stars. *MNRAS*, 409:481–490, December 2010. doi: 10.1111/j.1365-2966.2010.17329.x.
- [51] A. K. Kulkarni and M. M. Romanova. Variability Profiles of Millisecond X-Ray Pulsars: Results of Pseudo-Newtonian Three-dimensional Magnetohydrodynamic Simulations. *ApJ*, 633:349–357, November 2005. doi: 10.1086/444489.
- [52] H. Bradt. *Astrophysics Processes: The Physics of Astronomical Phenomena*. Cambridge University Press, 1st edition, 2008.
- [53] K. Viironen and J. Poutanen. Light curves and polarization of accretion- and nuclear-powered millisecond pulsars. *A&A*, 426:985–997, November 2004. doi: 10.1051/0004-6361:20041084.
- [54] M. C. Miller and F. K. Lamb. Bounds on the Compactness of Neutron Stars from Brightness Oscillations during X-Ray Bursts. *ApJL*, 499:L37, May 1998. doi: 10.1086/311335.
- [55] J. Poutanen and M. Gierliński. On the nature of the X-ray emission from the accreting millisecond pulsar SAX J1808.4-3658. *MNRAS*, 343:1301–1311, August 2003. doi: 10.1046/j.1365-8711.2003.06773.x.
- [56] A. M. Beloborodov. Gravitational Bending of Light Near Compact Objects. *ApJL*, 566:L85–L88, February 2002. doi: 10.1086/339511.
- [57] C. Cadeau, D. A. Leahy, and S. M. Morsink. Pulse Shapes from Rapidly Rotating Neutron Stars: Equatorial Photon Orbits. *ApJ*, 618:451–462, January 2005. doi: 10.1086/425857.
- [58] C. Cadeau, S. M. Morsink, D. A. Leahy, and S. S. Campbell. Light Curves for Rapidly Rotating Neutron Stars. *ApJ*, 654:458–469, January 2007. doi: 10.1086/509103.

- [59] S. M. Morsink, D. A. Leahy, C. Cadeau, and J. Braga. The Oblate Schwarzschild Approximation for Light Curves of Rapidly Rotating Neutron Stars. *ApJ*, 663:1244–1251, July 2007. doi: 10.1086/518648.
- [60] F. K. Lamb, S. Boutloukos, S. Van Wassenhove, R. T. Chamberlain, K. H. Lo, A. Clare, W. Yu, and M. C. Miller. A Model for the Waveform Behavior of Accreting Millisecond X-Ray Pulsars: Nearly Aligned Magnetic Fields and Moving Emission Regions. *ApJ*, 706:417–435, November 2009. doi: 10.1088/0004-637X/706/1/417.
- [61] NASA. RXTE Guest Observer Facility. URL <http://heasarc.gsfc.nasa.gov/docs/xte/>.
- [62] K. Jahoda, J. H. Swank, A. B. Giles, M. J. Stark, T. Strohmayer, W. Zhang, and E. H. Morgan. In-orbit performance and calibration of the Rossi X-ray Timing Explorer (RXTE) Proportional Counter Array (PCA). In O. H. Siegmund and M. A. Gummin, editors, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 2808 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pages 59–70, October 1996.
- [63] R. Wijnands. An observational review of accretion-driven millisecond X-ray pulsars. *NuPhS*, 132:496–505, June 2004. doi: 10.1016/j.nuclphysbps.2004.04.084.
- [64] E. Rivers, A. Markowitz, and R. Rothschild. Spectral Survey of X-ray Bright Active Galactic Nuclei from the Rossi X-ray Timing Explorer. *ApJS*, 193:3, March 2011. doi: 10.1088/0067-0049/193/1/3.
- [65] N. Shaposhnikov, K. Jahoda, C. Markwardt, J. Swank, and T. Strohmayer. Advances in the RXTE Proportional Counter Array Calibration: Nearing the Statistical Limit. *ApJ*, 757:159, October 2012. doi: 10.1088/0004-637X/757/2/159.

- [66] Indian Space Research Organisation. ASTROSAT: A satellite mission for multi-wavelength astronomy. URL <http://astrosat.iucaa.in/>.
- [67] K. Gendreau. NICER science overview. Technical report, NASA, 2013. URL http://heasarc.gsfc.nasa.gov/docs/nicer/slides/NICER.Science_Overview.pdf.
- [68] M. Feroci and LOFT Collaboration. The Large Observatory for X-ray Timing (LOFT). *Exp. Astron.*, 34:415–444, October 2012. doi: 10.1007/s10686-011-9237-2.
- [69] N. Stergioulas and J. L. Friedman. Comparing models of rapidly rotating relativistic stars constructed by two numerical methods. *ApJ*, 444:306–311, May 1995. doi: 10.1086/175605.
- [70] C. W. Misner, K. S. Thorne, and J. A. Wheeler. *Gravitation*. W. H. Freeman, 1973.
- [71] W. Rindler. *Essential Relativity: Special, General, and Cosmological*. Texts and Monographs in Physics. Springer-Verlag, 2nd edition, 1977.
- [72] M. R. Spiegel. *Mathematical Handbook of Formulas and Tables*. Schaum’s Outline Series. McGraw-Hill Publishing Co, 25th edition, 1990.
- [73] J. Terrell. Invisibility of the Lorentz Contraction. *Phys. Rev.*, 116:1041–1045, November 1959. doi: 10.1103/PhysRev.116.1041.
- [74] S. Chandrasekhar. *Radiative Transfer*. Dover Publications, Inc., 1960.
- [75] E. Kuulkers, P. R. den Hartog, J. J. M. in’t Zand, F. W. M. Verbunt, W. E. Harris, and M. Cocchi. Photospheric radius expansion X-ray bursts as standard candles. *A&A*, 399:663–680, February 2003. doi: 10.1051/0004-6361:20021781.

- [76] T. M. Braje, R. W. Romani, and K. P. Rauch. Light Curves of Rapidly Rotating Neutron Stars. *ApJ*, 531:447–452, March 2000. doi: 10.1086/308448.
- [77] Smithsonian Institution Chandra X-ray Center. Chandra X-ray Observatory: Energy bands, 2009. URL <http://cxc.harvard.edu/csc/columns/ebands.html>.
- [78] XMM catalogue public pages, 2008. URL http://xmmssc-www.star.le.ac.uk/Catalogue/xcat_public_2XMM.html.
- [79] Instruments on-board LOFT, 2012. URL <http://www.isdc.unige.ch/loft/index.php/instruments-on-board-loft>.
- [80] The LAXPC instrument, 2008. URL http://www.iucaa.ernet.in/~astrosat/laxpc_specs.html.
- [81] C. Darwin. *On the Origin of Species: By Means of Natural Selection*. Dover Thrift Editions, 2006.
- [82] D. E. Goldberg. *The Design of Innovation: Lessons from and for Competent Genetic Algorithms*. Kluwer Academic Publishers, 2002.
- [83] D. E. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley Publishing Company, Inc., 1989.
- [84] J. D. Fiege. *Qubist User's Guide: Optimization, Data-Modeling, and Visualization with the Qubist Global Optimization Toolbox for MATLAB*. nQube Technical Computing Corporation, 4th edition, March 2010.
- [85] A. Rogers and J. D. Fiege. Gravitational Lens Modeling with Genetic Algorithms and Particle Swarm Optimizers. *ApJ*, 727:80, February 2011. doi: 10.1088/0004-637X/727/2/80.

- [86] D. Grant. Phys 699 class lectures [*unpublished*], Fall 2013.
- [87] A. W. Steiner, J. M. Lattimer, and E. F. Brown. The Equation of State from Observed Masses and Radii of Neutron Stars. *ApJ*, 722:33–54, October 2010. doi: 10.1088/0004-637X/722/1/33.
- [88] Z. Wang, R. P. Breton, C. O. Heinke, C. J. Deloye, and J. Zhong. Multi-band Studies of the Optical Periodic Modulation in the X-Ray Binary SAX J1808.4-3658 during Its Quiescence and 2008 Outburst. *ApJ*, 765:151, March 2013. doi: 10.1088/0004-637X/765/2/151.
- [89] D. A. Leahy, S. M. Morsink, and C. Cadeau. Limits on Mass and Radius for the Millisecond-Period X-Ray Pulsar SAX J1808.4-3658. *ApJ*, 672:1119–1126, January 2008. doi: 10.1086/523794.
- [90] F. Özel. Surface emission from neutron stars and implications for the physics of their interiors. *Rep. Prog. Phys.*, 76(1):016901, January 2013. doi: 10.1088/0034-4885/76/1/016901.