

Radiation Drag on the Accretion Disc in Type-I X-ray Bursts

A thesis submitted for the degree of

Doctor of Philosophy

by

Hauke Worpel

B.Sc., University of Melbourne

Grad.Dip.Sc.(Hons), Monash University

Department of Mathematical Sciences

Monash University

Australia

January 2015

Contents

	vii
Declaration	ix
1 Introduction	1
1.1 Type-I X-ray bursts	3
1.1.1 Burning regimes	8
1.2 Radius expansion bursts	11
1.3 Mass and radius measurements from bursts	12
1.3.1 Gravitational redshift	12
1.3.2 Cooling tails	15
1.3.3 Radius Expansion Bursters as Standard Candles	16
1.4 Data analysis	16
1.4.1 Spectral shape changes	18
1.5 Computer simulations	20
1.6 Accretion disc structure	23
1.6.1 Vertical disc structure	27
1.6.2 Previous theoretical analyses	28
1.7 Outline of this thesis	29
2 Methods and Data Analysis	31
2.1 Data collection and reduction	31
2.1.1 Introduction to X-ray spectral analysis	31
2.1.2 Burst spectral analysis with <i>RXTE</i>	33
2.2 Method	35

2.2.1	Characterising the persistent emission	36
2.2.2	Modelling the burst spectrum	37
2.2.3	Burst selection	43
2.2.4	Modelling the persistent emission	43
2.2.5	Fitting burst spectra	46
2.3	Eddington Fluxes	47
3	SPH Modelling	51
3.1	Fluid dynamics	51
3.1.1	Conservation of mass	51
3.1.2	Conservation of momentum	52
3.1.3	Conservation of Energy	53
3.2	Formulation in Smoothed Particle Hydrodynamics	55
3.2.1	Overview of SPH	55
3.2.2	Kernel choice	56
3.2.3	Smoothing length	57
3.2.4	Error estimates	59
3.3	Fluid equations in SPH	60
3.3.1	Conservation of mass in SPH	60
3.3.2	Conservation of momentum in SPH	61
3.3.3	Timestepping in PHANTOM	62
3.4	The <code>extern_prdrag</code> module	63
3.4.1	Description of Poynting-Robertson drag parameter β	66
4	Radius Expansion Bursts	69
4.1	Distribution of f_a	69
4.2	f_a as a function of burst stage	72
4.3	Spectral shape changes	75
4.3.1	Phenomenological modification to burst spectrum	75
4.3.2	Does a modification to the burst spectrum remove the need for variable persistent emission?	77
4.3.3	Correlation with burst flux?	77

4.3.4	Neutron star atmosphere spectra	78
4.4	Comparison with theory	79
4.5	Luminosity of bursters in globular clusters	81
4.5.1	Data	83
4.5.2	Statistical method	84
4.5.3	Priors	85
4.5.4	Results	85
5	Non-Eddington Limited Bursts	89
5.1	Data analysis	89
5.1.1	Tests of significance	92
5.1.2	Correlation with burst flux	94
5.1.3	Collation of Eddington fluxes	96
5.2	Is the persistent emission spectral shape constant?	97
5.2.1	Short-term variability	97
5.2.2	Medium-term variability	98
5.2.3	Detectability of persistent emission spectral shape change	99
5.2.4	Variability outside the 2.5-20 keV energy range	100
5.3	Eddington fluxes from non-PRE bursts	102
5.3.1	Estimating γf_a upper limit	103
5.3.1.1	Application to sources with known Eddington flux	105
5.3.1.2	Application to sources with unknown Eddington flux	106
6	SPH simulation setup and results	109
6.1	Simulation Setup	110
6.1.1	Neutron star	110
6.1.2	Disc	110
6.2	Numerical tests	112
6.2.1	Correct integration of equations of motion	112
6.2.2	Resolution of scale height	113
6.2.3	Convergence of accretion rate	114
6.3	Burst profile	115

6.4	Prescriptions for β	117
6.4.1	Time-independent β	117
6.4.2	Alternative approaches	119
6.4.2.1	Approach A	119
6.4.2.2	Approach B	119
6.4.2.3	Approach C	121
6.4.2.4	Approach D	121
6.4.3	Feedback Luminosity	122
6.4.3.1	Accretion of individual SPH particles	123
6.5	Results	123
6.5.1	Full simulations	123
6.5.2	Low resolution simulations	124
6.6	Discussion	125
7	Conclusions and Future Work	143
7.1	Discussion	143
7.1.1	Structure of the accretion disc	145
7.2	Difficulties	147
7.2.1	Degeneracy of spectral components	147
7.2.2	Magnitude of accretion enhancement	148
7.3	Future prospects	148
7.3.1	Theory	148
7.3.2	Observational work	149
	Appendix A	151
	Appendix B	155

To my family and friends.

Declaration

In accordance with Monash University Doctorate Regulation 17.2 Doctor of Philosophy and Research Masters regulations the following declarations are made:

I hereby declare that this thesis contains no material which has been accepted for the award of any other degree or diploma at any university or equivalent institution and that, to the best of my knowledge and belief, this thesis contains no material previously published or written by another person, except where due reference is made in the text of the thesis.

This thesis includes one original paper published in peer reviewed journals ([Worpel et al. 2013](#)) and one accepted publication ([Worpel et al. 2015](#)), partially rewritten for this thesis. The core theme of the thesis is to investigate the effect of burst luminosity on the accretion flow in type I X-ray bursts. The ideas, development and writing up of all the papers in the thesis were the principal responsibility of myself, the candidate, working within the Monash Centre for Astrophysics under the supervision of Dr. Duncan Galloway and Dr. Daniel Price.

Under the Copyright Act 1968, this thesis must be used only under the normal conditions of scholarly fair dealing. In particular no results or conclusions should be extracted from it, nor should it be copied or closely paraphrased in whole or in part without the written consent of the author. Proper written acknowledgement should be made for any assistance obtained from this thesis.

I certify that I have made all reasonable efforts to secure copyright permissions for third-party content included in this thesis and have not knowingly added copyright content to my work without the owner's permission.

Table 1. Contributions to Published Work

Chapter	Publication Title	Publication Status	Contribution
2, 4	Evidence for Accretion Rate Change During Type-I X-ray Bursts	Published	
5	Evidence for enhanced persistent emission during sub-Eddington thermonuclear bursts	Accepted	

I have renumbered sections of submitted or published papers in order to generate a consistent presentation within the thesis.



Signed

22 January 2015

Date

Summary

Type I X-ray bursts on accreting neutron stars offer a means to probe the interior of these stars by providing information about the conditions at its surface, via the measurement of radius and surface gravitational redshift. Our understanding has been hampered by difficulties in finding spectral models that adequately describe the spectra, in part because we do not properly understand how the burst luminosity influences the structure of the disc and the properties of the accretion flow.

I have used archival *Rossi X-ray Timing Explorer* data to perform spectral analyses of 1,759 bursts from 56 sources. I investigated the effect of allowing the pre-burst persistent emission to vary with time, while holding its spectral shape fixed. I found that an increase in the intensity, to several times its pre-burst level, significantly improves the quality of the spectral fits. This increase cannot be attributed to changes in the shape of either the burst or the persistent components. One possible interpretation of these results is a temporary increase in accretion rate, possibly as a consequence of radiation drag on the disc. I have shown that the magnitude of the persistent flux increase is consistent with the results of earlier computer simulations.

I further test this hypothesis using smoothed particle hydrodynamics (SPH) computer simulations of an irradiated accretion disc. The simulations evolve the Navier-Stokes equations with an extra velocity-dependent force representing radiation drag. I test the effect of partial self-shading of the disc by devising four methods for propagating radiation within it. The results suggest that large accretion enhancements are possible provided that angular momentum loss can be efficiently communicated to the interior of the disc.

Acknowledgements

The author has been supported by an APA postgraduate research scholarship. This thesis utilizes preliminary analysis results from the Multi-INstrument Burst ARchive (MINBAR)¹, which is supported under the Australian Academy of Science’s Scientific Visits to Europe program, and the Australian Research Council’s Discovery Projects and Future Fellowship funding schemes. This research has made use of data obtained through the High Energy Astrophysics Science Archive Research Center Online Service, provided by the NASA/Goddard Space Flight Center. Some of the computational work was performed on the gSTAR national facility at Swinburne University of Technology. gSTAR is funded by Swinburne and the Australian Governments Education Investment Fund. Some of the computational work was performed on Monash Sun Grid.

The author is grateful to the International Space Science Institute (ISSI) in Bern for the support of an International Team on Type I X-ray Bursts, especially for compiling much of the data in Table B.2. I am also grateful to the anonymous referees on [WGP13](#) and [Worpel et al. \(2015\)](#).

Finally, I am thankful to my supervisors, Daniel Price and Duncan Galloway, for their constant encouragement and helpful advice.

¹see burst.sci.monash.edu/minbar

This work has made use of the following computer tools and utilities:

- Astropy ([Astropy Collaboration et al. 2013](#))
- FTOOLS ([Blackburn 1995](#))
- Maxima ([Maxima 2012](#))
- Pari-GP ([The PARI Group 2011](#))
- SPLASH ([Price 2007](#))
- TOPCAT ([Taylor 2005](#)) and STILTS ([Taylor 2006](#))
- XSPEC ([Arnaud 1996](#))

Chapter 1

Introduction

Low-mass X-ray Binaries

A low-mass X-ray binary (LMXB) is a binary star system in which the primary star is a compact object (either a neutron star or a black hole) and the companion is a main-sequence or evolved star, with a mass comparable to that of the Sun, that transfers material to the other star through Roche lobe overflow.

Since material taken from the companion must lose angular momentum to arrive on the primary, an accretion disc forms. The accretion process is bright in X-rays, emitting in excess of 10^{34} erg s⁻¹ and up to 10^{37} erg s⁻¹. It is estimated that there are a few hundred LMXB systems in the Galaxy ([Rosswog & Brüggen 2007](#)).

The LMXB system must remain intact after the supernova explosion that produces the compact object, and it must also end up with an orbital separation small enough for the companion to fill its Roche lobe. Since the companion is a low-mass star its Roche lobe is small, so the orbital separation must also be small. Orbital periods for LMXBs are less than ten days, and can be as little as a few tens of minutes for ultracompact X-ray binaries (UCXB) where the companion is a white dwarf ([Podsiadlowski et al. 2002](#)). In many cases the primary must once have been a giant star with a physical size larger than the current separation of the system, so some mechanism for subsequently shrinking the orbit must

exist. These stringent requirements account for the rarity of LMXBs (Rosswog & Brüggen 2007).

Because the companion is long-lived, the primary can accrete for a long time (10^6 to 10^9 years). LMXBs are therefore an old population that have had plenty of time to move away from where they were formed. LMXBs are concentrated towards the Galactic center and somewhat towards the Galactic plane (Lewin et al. 1977), as shown in Figure 1.1, but are not associated with regions of active star formation. LMXBs are preferentially found in globular clusters, and it may be that the gravitational influence of other stars in the globular cluster provides a means of decreasing the orbital separation (Podsiadlowski et al. 2002).

In principle a neutron star might accrete enough material directly from interstellar space to show comparable X-ray accretion emission and bursting behaviour to a star accreting from a companion. This is possible if it happens to be located inside a dense gas cloud (Ostriker et al. 1970). However, the high Galactic latitude of many bursters, away from the dense parts of the Galaxy, argues against accretion from interstellar space being a plausible mechanism for supplying the neutron star with gas (e.g., Lewin & Joss 1981). Nor is it likely that a neutron star could be left with a massive disc retained from its creation. A companion star is therefore necessary to provide the material.

LMXBs are conspicuous in X-rays but very dim ($M_V \sim 18$; Liu et al. 2001, 2007) at optical wavelengths, due to the small size of the luminous compact star, and because the companion is a dim star. Many LMXBs are persistent sources, meaning that they emit continuously. Others are transient, with very low ($10^{30} - 10^{32}$ erg s $^{-1}$) X-ray luminosity most of the time, and only periodic episodes of high accretion. Transience is thought to arise from a disc instability: a cold disc with low viscosity accumulates a large amount of gas, until it becomes warm enough for ionisation to occur. Then the viscosity rises quickly, causing the accretion rate to increase until the disc depletes enough to de-ionise it. Then the system returns to quiescence (e.g., Rosswog & Brüggen 2007). Some transient systems, like the Rapid Burster, alternate between outburst and quiescence every few months (e.g., Lewin et al. 1993, Masetti 2002). Others can remain in quiescence for decades, for example Cen X-4.

Many LMXBs exhibit very rapid increases in X-ray flux, to many times their previous levels, on timescales of seconds. These events are called X-ray bursts. They only occur in LMXBs where the primary is a neutron star, and do not occur when the system is in quiescence. X-ray bursts and their influence on the accretion disc are the subject of this thesis.

1.1 Type-I X-ray bursts

Thermonuclear (type I) X-ray bursts arise from the unstable ignition of accreted hydrogen and/or helium near the surface of an accreting neutron star in a LMXB system (Fujimoto et al. 1981, Strohmayer & Bildsten 2006). Stellar material accreted from a low-mass companion star accumulates on the surface of the neutron star, forming a layer about ten meters in thickness (Lewin & Joss 1981), where it is compressed and heated hydrostatically. When the temperature and pressure are high enough at the base of the accumulated material, approximately 2×10^9 K and 4×10^{22} erg cm⁻³ respectively (Strohmayer & Bildsten 2006), a thermonuclear explosion is triggered. These events are observed as a sudden increase in the X-ray luminosity to many times the pre-burst level (e.g., Lewin et al. 1993, Strohmayer & Bildsten 2006). Type I bursts were discovered in 1975 (Grindlay et al. 1976), when bursts from 4U 1820–30 were detected by the *Astronomische Nederlandse Satelliet* X-ray satellite. Thermonuclear bursts are the primary method for distinguishing LMXBs containing a neutron star as the compact companion from those containing a black hole, because neutron stars have a solid surface on which bursts can occur and black holes do not.

Typical bursts exhibit rise times from about 1 second to 10 seconds, durations of a few tens of seconds to a few minutes and have total energy outputs of $10^{39} - 10^{40}$ erg (e.g., Day & Tawara 1990, Gottwald et al. 1986). Temperatures during these bursts are of the order of a few keV ($\sim 10^7$ K), so the peak of the emission is in X-rays. These bursts are known as type I thermonuclear bursts, to distinguish them from type II bursts which are thought to be caused by sudden accretion events (Lewin et al. 1993). The peak luminosity can be as

high as 10^{38} erg s $^{-1}$, making them easily detectable throughout the Galaxy. Type I X-ray bursts have been detected from 104 different sources within the Galaxy¹.

The spectra of X-ray bursts are generally well described by a blackbody function (e.g., Swank et al. 1977):

$$B(E)dE = \frac{K E^3}{\exp[E/kT] - 1}, \quad (1.1.1)$$

where K is a variable normalisation parameter that depends on the size and distance of the emitting surface, and kT is the colour temperature, which is usually expressed in keV. The temperature typically has peak values of $kT \approx 2 - 3$ keV ($2 - 3.5 \times 10^7$ K, e.g., Galloway et al. 2008, hereafter G08). Time resolved spectral fitting shows evidence of an initial rise in temperature, followed by a slower decrease after the peak luminosity is reached. The decrease has usually been found to be approximately exponential, though theory suggests that a power law is more appropriate (in't Zand, Cumming, Triemstra, Mateijsen & Bagnoli 2014). The behaviour of kT is naturally interpreted as initial heating of the photosphere by the nuclear burning followed by slower cooling after the nuclear fuel is all consumed; see Figure 1.2.

It usually takes ~ 0.1 s for the burning region to spread over the entire surface of the star (Day & Tawara 1990). This is the timescale for *deflagration*, whereby burning is initiated in neighbouring material by radiative heat transfer. In some powerful bursts, which reach the Eddington limit in less than one millisecond (in't Zand, Keek & Cavecchi 2014), *detonation* may occur; here the spreading of the burning area is caused by a powerful, fast-moving shock wave which compresses material ahead of it (e.g., in't Zand, Cumming, Triemstra, Mateijsen & Bagnoli 2014) and is faster than deflagration by roughly three orders of magnitude.

Evidence that nuclear burning is responsible for bursts comes from considerations of the energetics. The fuel that burns in a burst was accreted onto the neutron star surface, liberating its gravitational potential energy as radiation. Every nucleon falling to the neutron star surface emits energy twice; once when it is accreted and once when it is burnt.

¹<http://www.sron.nl/~jeanz/bursterlist.html> and <http://burst.sci.monash.edu/wiki/index.php?n=MINBAR.SourceTable>

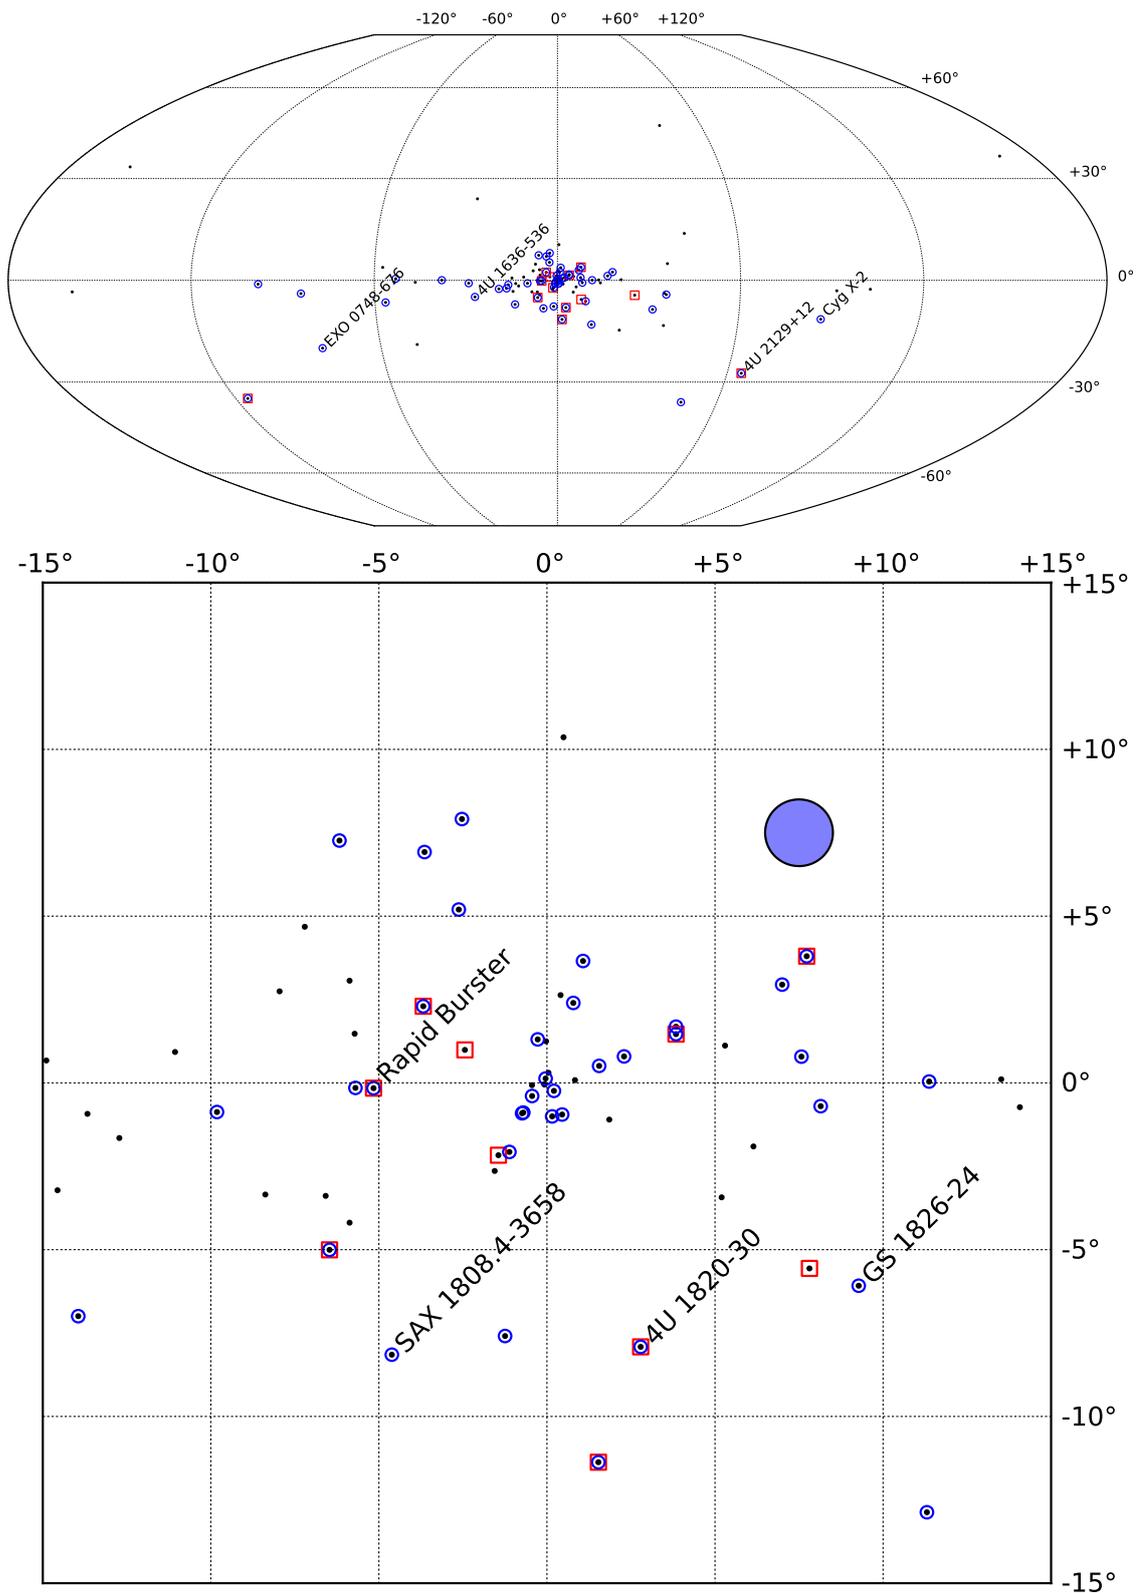


Figure 1.1 Locations of the known bursting sources in the Galaxy. The plot is in Galactic coordinates, with the Galactic center at the center of the plot. The upper panel shows the entire sky, and the lower panel shows the region around the Galactic center where the bursters are concentrated. The 1° field of view of *RXTE* is indicated as the blue circle in the lower panel. Sources studied in this thesis are indicated with blue circles, and those in globular clusters are indicated with a red square.

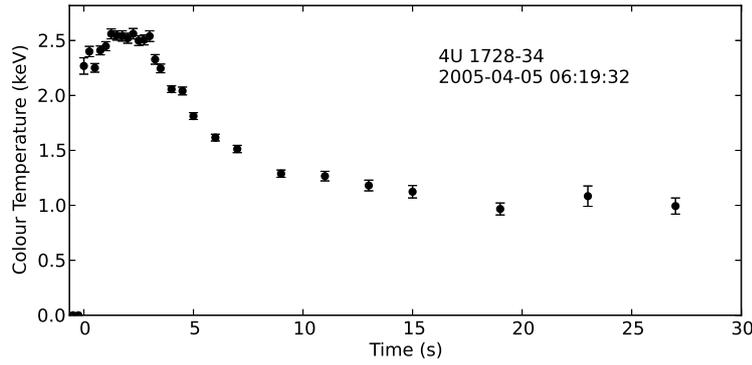


Figure 1.2 Temperature against time for a burst from the neutron star 4U 1728–34. It shows a rapid rise in temperature as the nuclear burning spreads across the surface of the star and heats the atmosphere, followed by more gradual cooling after all the nuclear fuel has been consumed.

The luminosity, or rate of energy release, obtained by accretion is

$$L_{\text{Acc}} = \frac{dE}{dt} = \dot{M} \frac{GM_*}{R_*}, \quad (1.1.2)$$

where M_* and R_* are the mass and radius of the neutron star and \dot{M} is the accretion rate. This is equivalent to about 180 MeV/nucleon for the “canonical” neutron star of mass $M_* = 1.4M_\odot$ and radius $R_* = 10$ km. The bulk of this energy is liberated as X-rays, and so the X-ray luminosity is a proxy for accretion rate (Verbunt et al. 1984). By comparison, the amount of energy released by nuclear burning is only 7 MeV per nucleon for the burning of hydrogen. The ratio of these, α , is of the order of 25 to 100, depending on the composition of the fuel and the mass and radius of the star (Lewin et al. 1993), with lower values of α indicating ignition in mixed H/He fuel. The quantity α can also be expressed as

$$\alpha = \frac{E_p}{E_b} = \frac{\Delta t F_p}{E_b}, \quad (1.1.3)$$

where E_b is the burst fluence, and E_p and F_p are the flux and fluence of persistent emission between bursts. These quantities can be measured from observations of bursting sources long enough to include multiple bursts, and the α values measured from observations are consistent with nuclear burning. For example, the prolific burster 4U 1636–536 typically shows α values in the range 40–100 (G08), and the helium accreting burster 4U 1820–30 has higher α , ranging from 125 to 155 (Haberl et al. 1987).

That the majority of the accreting material's gravitational potential energy is released near the stellar surface can be shown with a simple argument. A particle of mass m descending from a radius R to $R - \Delta R$ has its gravitational energy change by

$$\Delta E_G = \frac{-GM_* m \Delta R}{R(R - \Delta R)}. \quad (1.1.4)$$

As $\Delta R \rightarrow 0$ this becomes

$$\frac{dE_G}{dR} = -\frac{GM_* m}{R^2}, \quad (1.1.5)$$

thus the rate of potential energy loss increases with decreasing altitude. If the particle remains approximately in a circular Keplerian orbit, then it gains kinetic energy and it is easy to show that

$$\frac{dE_K}{dR} = \frac{1}{2} \frac{GM_* m}{R^2} = \frac{1}{2} \frac{dE_G}{dR}. \quad (1.1.6)$$

That is, half of its gravitational potential energy is radiated away while the other half is converted to kinetic energy. Furthermore, when the particle is accreted onto the stellar surface it must come to rest. Some of its kinetic energy goes into spinning up the star, and the rest is radiated away.

The braking of material to match the rotation velocity of the star is thought to be a turbulent process, and may contribute to the heating of layers deeper in the atmosphere (Inogamov & Sunyaev 2010). These bursting systems are strongly believed to have low magnetic field ($\lesssim 10^{12}$ G, e.g., Treves et al. 1988, and less than 10^{10} G for most bursting systems), because if the magnetic field were high, the accreting material would be funnelled onto the magnetic poles of the star where, due to the very high *local* accretion rate, it would burn stably and not burst (Joss 1978; see also §1.1.1). Furthermore, the rapid rotation of the neutron stars in LMXBs implies a low magnetic field strength. Some authors argue that the magnetic field of neutron stars decays over time (e.g., Alpar et al. 1981), and it is also possible that the accreting material buries the magnetic field in a LMXB (Cumming et al. 2001).

When the accretion rate from the companion is low, below about 1% \dot{M}_{Edd} , the disc is truncated at some radius away from the neutron star surface. Accreting material passing through this radius instead becomes a spherical accretion flow, and it has long been

understood (Alme & Wilson 1973) that spherical accretion onto a neutron star produces a spectrum that closely resembles a blackbody. When the accretion rate is high, the disc presumably extends to the neutron star surface. The X-ray spectrum associated with discs of this kind tend to be “harder”, that is, they produce more photons at high energies than the “soft” blackbody spectrum.

1.1.1 Burning regimes

The global accretion rate determines the behaviour of bursts, and can be divided into several classes of behaviour. At the lowest accretion rates ($< 1\% \dot{M}_{\text{Edd}}$), the gas accumulates at low temperature until the hydrogen ignites and burns via the hot (β -limited) CNO cycle (see Figure 1.4). Though hydrogen burning on its own is not fast enough to cause a burst, because its rate is limited by the β decays, it immediately triggers triple-alpha helium burning (see Figure 1.5) in a hydrogen rich environment.

At moderate accretion rates, between about 1% and 10% \dot{M}_{Edd} , the temperature is sufficient for stable hydrogen burning, and the hydrogen that is accreted is burnt as quickly as it arrives. The bursts occurring in this case are unstable helium burning, via the triple- α process (see Figure 1.5), in a hydrogen poor environment. These have rapid rise times, of less than one second (see Figure 1.3) because the triple- α reaction rates increase with increasing temperature and there are no β decays to slow the process.

When both hydrogen and helium are present, the nuclei of the hot CNO cycle can absorb a helium-4 nucleus in reactions such as $^{15}\text{O}(\alpha, \gamma)^{19}\text{Ne}$ and $^{14}\text{O}(\alpha, p)^{17}\text{F}$. These “breakout” nuclei open further pathways for both hydrogen and helium to be burnt, resulting in the production of a large variety of heavier nuclei proceeding as high as iron. The burning rate is moderated partially by β decays of proton-rich nuclei. Such bursts are characterised by a relatively slow rise time, typically a few seconds. Helium bursts in a hydrogen rich environment occur when the accretion rate is high enough for hydrogen to accumulate faster than it can be burnt.

It is also believed that such hydrogen-rich bursts can occur at very low accretion rates, below about 1% \dot{M}_{Edd} . In this case the hydrogen does not get hot enough to burn stably and only begins to burn, unstably, once a sufficient quantity of hydrogen has accreted.

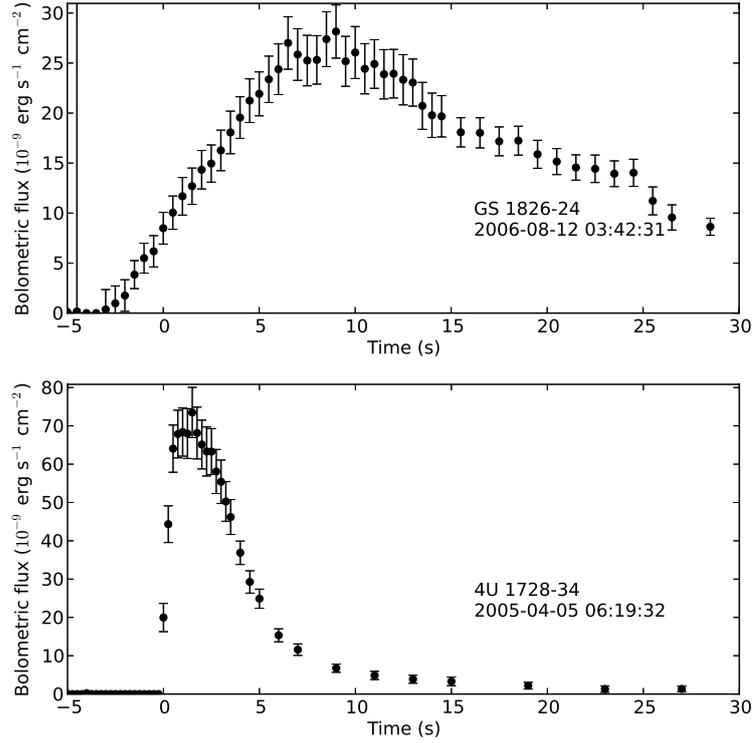


Figure 1.3 Bolometric flux, calculated from blackbody temperature and normalisation, against time for two bursts occurring in different atmospheric compositions. The top panel shows a burst occurring in a hydrogen rich environment. It has a slow rise time. The bottom panel shows a burst occurring in a hydrogen-poor atmosphere. Here the rise time is very much shorter.

Once the hydrogen begins to burn, a helium burst is immediately triggered in the still hydrogen-rich atmosphere. No such bursts have yet been observed.

Naturally, if the companion star is a white dwarf, there is no hydrogen and so all the bursts are the rapidly rising events regardless of accretion rate.

The rise times and durations of bursts can therefore be used to infer the composition of the atmosphere. Naturally these distinctions are only relevant if the accreted material contains a significant amount of hydrogen. If the donor is a white dwarf, as in the case of 4U 1820–30 (King & Watson 1986) and probably also 4U 1728–34 (Galloway et al. 2010b), all bursts are the rapidly rising hydrogen poor bursts. The divisions between the regimes in hydrogen accretors depends upon the proportion of hydrogen in the gas, the metallicity (i.e., the amount of CNO catalyst), and the temperature of the neutron star crust. All these factors influence the range of accretion rates where stable hydrogen burning occurs.

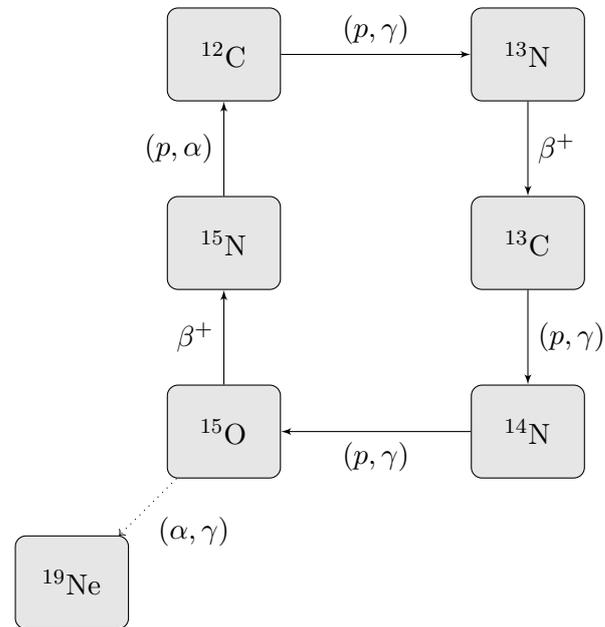


Figure 1.4 The hot (i.e. β -decay limited) CNO cycle. It takes the form of a cycle in which an initial carbon-12 nucleus absorbs four protons and, through two beta decays and the emission of a helium-4 nucleus, returns to carbon-12. Also shown is the breakout reaction $^{15}\text{O}(\alpha, \gamma)^{19}\text{Ne}$, which removes CNO catalyst from the cycle.

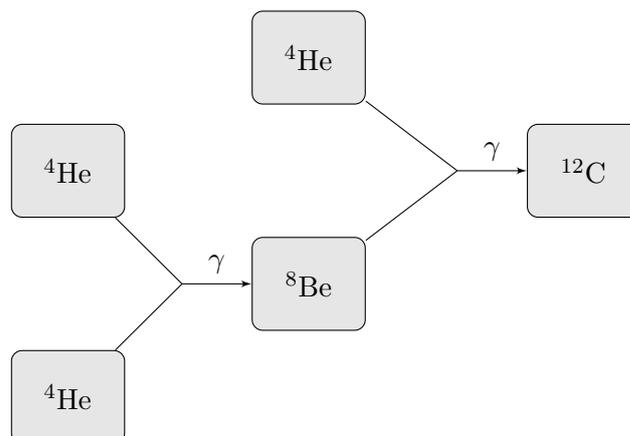


Figure 1.5 The triple- α process, in which three helium-4 nuclei combine to form one nucleus of carbon-12.

1.2 Radius expansion bursts

During an X-ray burst the liberated energy from nuclear burning goes into heating the gas; this causes both the luminosity and temperature of the observed stellar spectrum to increase. For a certain fraction of bursts, the luminosity eventually reaches the Eddington luminosity, the luminosity at which outwards radiation pressure balances the inward pull of gravity. This limit is derived as follows:

Consider a quantity of gas with mass m located at a distance R from a star of mass M and luminosity L . The force of gravity on the gas is

$$F_{grav} = \frac{GMm}{R^2}, \quad (1.2.1)$$

and the radiation pressure at radius R is

$$P_{rad} = \frac{L}{c} \frac{1}{4\pi R^2}, \quad (1.2.2)$$

where G and c are the gravitational constant and the speed of light respectively. Radiation pressure is related to radiation force through

$$F_{rad} = P_{rad}\kappa m, \quad (1.2.3)$$

where κ is the opacity (cross-sectional area per unit mass) of the gas. When the forces of gravity and radiation pressure are equal we have

$$\frac{GMm}{R^2} = \frac{L\kappa m}{4\pi c R^2}, \quad (1.2.4)$$

from which it follows that the Eddington luminosity

$$L_{Edd} = \frac{4\pi GMc}{\kappa}. \quad (1.2.5)$$

When the Eddington luminosity is reached, the excess energy is efficiently converted to kinetic energy of the atmosphere and the luminosity is thought to remain within a few percent of the local Eddington luminosity throughout the radius expansion (e.g., [Hanawa &](#)

Sugimoto 1982, Ebisuzaki et al. 1983, Titarchuk 1994b). Such events are called *photospheric radius expansion* (PRE) bursts. After the end of the nuclear burning the atmosphere then expands and cools, until the burst flux drops below the Eddington limit again. Then the atmosphere returns to the stellar surface.

1.3 Mass and radius measurements from bursts

Bursts have long been recognised as a means of potentially measuring the mass and radius of neutron stars, and thereby deriving the equation of state (e.g., van Paradijs 1979, Damen et al. 1990, Lattimer & Prakash 2007). This is because bursts provide a probe into conditions at the neutron star surface. Attempts to calculate these fundamental properties include calculating the surface gravitational redshift (i.e., the ratio of mass to radius) from PRE bursts, and measuring the neutron star radius from the thermal emission of the photosphere as it is cooling after a burst.

1.3.1 Gravitational redshift

Although the luminosity of a neutron star undergoing a radius expansion burst remains at the Eddington luminosity throughout the radius expansion episode, the Eddington luminosity itself as measured by a distant observer depends upon the gravitational redshift factor and therefore upon the radius of the emitting photosphere. Observations of radius expansion bursts in principle allow the gravitational redshift factor at the neutron star's surface to be inferred, which means the radius and mass of the neutron star can be calculated (Damen et al. 1990).

The Eddington luminosity as measured by a distant observer $L_{edd,\infty}$ is not equal to the Eddington luminosity as measured by an observer travelling with the expanding atmosphere, with effective radius R , $L_{edd,R}$ due to gravitational redshift. The two quantities are described by the relation

$$L_{edd,\infty} = (1 + z_R)L_{edd,R} = L_{edd,R}\sqrt{1 - \frac{2GM}{Rc^2}}. \quad (1.3.1)$$

Thus, radius expansion bursts can in principle be used to measure the surface gravitational redshift of the neutron star and thereby its compactness (e.g., [Damen et al. 1990](#)), giving insights into the neutron star equation of state ([Lattimer & Prakash 2007](#)). This in turn may have implications for the structure of the disc, since for some stiff equations of state the radius of the innermost stable orbit is actually inside the neutron star ([Kluźniak & Wagoner 1985](#)).

There is a further complication arising, from the influence of temperature on opacity. I begin with the expression of the Eddington luminosity, as observed by a distant observer, as a function of temperature:

$$L_{Edd,\infty}(T) = \frac{4\pi cGM_*}{\kappa(T)} \left[1 - \frac{2GM_*}{c^2 R(T)} \right]^{1/2}, \quad (1.3.2)$$

where $\kappa(T)$ is the opacity of the material and all the other symbols have their usual meanings.

Let us now consider the temperature dependence of R . If the photosphere is radiating as a blackbody these two quantities are related by

$$T^4 R^2 = B^2 \quad (1.3.3)$$

for some constant B , and so $R = BT^{-2}$.

Continuing, the electron scattering coefficient κ is well-approximated over the temperature ranges under consideration by the formula

$$\kappa(T) \approx 0.2(1 + X) \left[1 + \left(\frac{kT}{39.2\text{keV}} \right)^{0.86} \right]^{-1}, \quad (1.3.4)$$

where X is the mass fraction of hydrogen in the atmosphere ([Lewin et al. 1993](#); see their equation 4.13b). The approximation is accurate to the tabulated values of [Buchler & Yueh \(1976\)](#) to about 4% ([Paczynski 1983b](#)).

Combining these, we have

$$L_{Edd,\infty}(T) = \frac{4\pi cGM_*}{0.2(1 + X)} \left[1 + \left(\frac{kT}{39.2\text{keV}} \right)^{0.86} \right] \left(1 - \frac{2GM_* T^2}{c^2 B} \right)^{1/2} \quad (1.3.5)$$

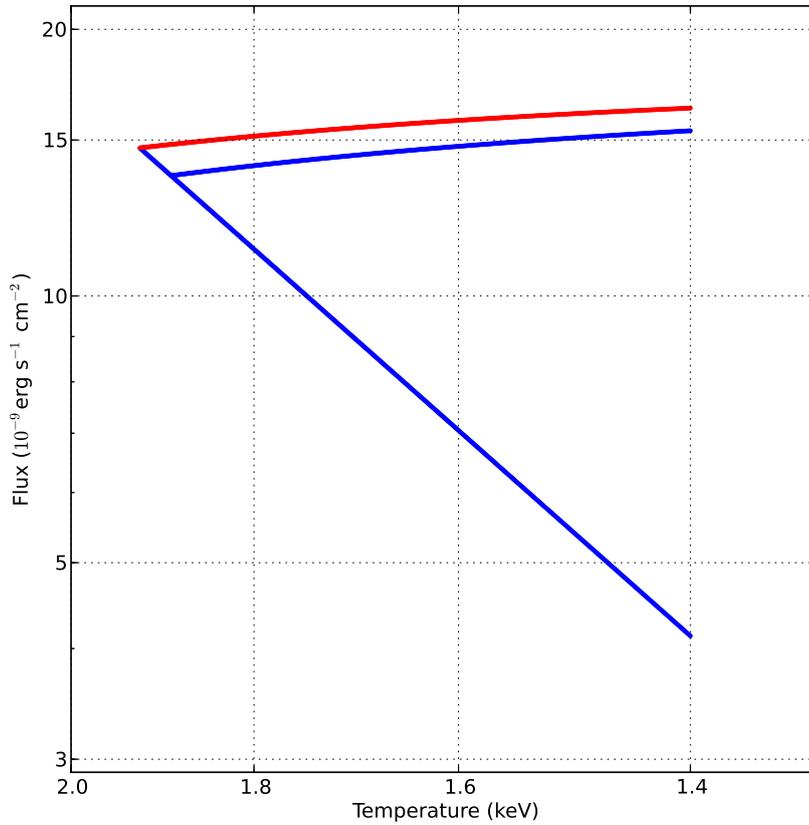


Figure 1.6 Flux-temperature curves for a theoretical PRE burst on a neutron star of 1.4 solar masses and 10 km radius, located 10 kpc away and occurring in an atmosphere of 70% hydrogen by weight. The red curve represents the burst when the temperature dependence on opacity is taken into account; the blue curve does not. It is clear that the temperature dependence on opacity diminishes the effect of gravitational redshift, reducing the potential to measure the surface gravitational redshift. The straight diagonal segment represents the *cooling tail*, when the atmosphere has returned to the surface and cools as a blackbody of constant radius.

It follows that the higher the temperature, the higher the Eddington luminosity. This effect partially suppresses the gravitational redshift effect on Eddington luminosity and reduces this method's ability to gauge the surface gravitational redshift (see Figure 1.6). Learning how to deal with this effect and others is crucial to understanding these systems. Furthermore, the flux-temperature curves of real bursts do not generally resemble the clean theoretical behaviour, as shown in Figure 1.7.

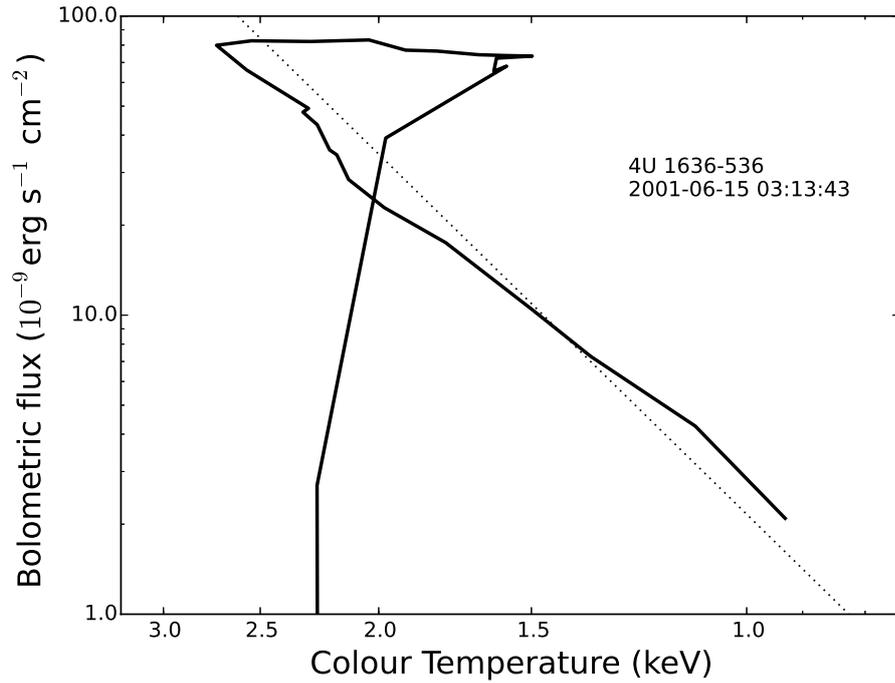


Figure 1.7 Flux-temperature curve for a typical radius expansion burst. The burst start is at the bottom left. The burst clearly does not show the orderly behaviour of the theoretical burst shown in Figure 1.6. Also shown is a line of constant radius (dotted line), 8.5 km at a distance of 6 kpc.

1.3.2 Cooling tails

A related approach, for measuring the radius of the neutron star, is to follow its flux-temperature curve down the cooling tail. In this method it is assumed that the entire surface of the star is emitting as a blackbody, which therefore should show decreasing temperature at constant radius (see the straight diagonal segment in Figure 1.6). This approach has been extensively studied (e.g., [Özel 2006](#), [Steiner et al. 2010](#), [Güver et al. 2012](#), [Özel et al. 2012](#)) but accurate radius determinations have been difficult to obtain, in part because the method depends on knowing accurately the distance to the neutron star, and in part because the cooling tails generally do not follow lines of constant radius (see Figure 1.7).

It is clear that attempts to determine the mass and radius of a neutron star from either radius expansion episodes or from the cooling tail are hampered by our lack of understanding of the physical behaviour of these systems and our interpretation of their spectra.

1.3.3 Radius Expansion Bursters as Standard Candles

The greatest uncertainty in the Eddington luminosity of a PRE-bursting neutron star is the composition of the expanding photosphere. The limiting flux is a factor of 1.7 lower for an atmosphere of cosmic composition than one which is composed of pure helium (e.g., [Galloway et al. 2008](#)). Despite this, the Eddington luminosities of many bursting neutron stars fall within a narrow range. For PRE-bursting neutron stars located in globular clusters, and whose distances can therefore be independently measured, about two thirds of the stars have maximum Eddington luminosities in the range $3.79 \pm 0.15 \times 10^{38}$ erg/s ([Kuulkers et al. 2003](#)).

PRE-bursts can be used as standard candles for determining distances, provided that individual sources attain the same limiting luminosity from burst to burst and that some means can be found for distinguishing those bursters that fall within the narrow luminosity range measured by [Kuulkers et al. \(2003\)](#) from those that do not. These standard candles will be relevant on distances of the order of the size of the Galaxy.

1.4 Data analysis

Spectral analyses of type I bursts typically make a number of implicit assumptions, namely (e.g., [Swank et al. 1977](#), [Kuulkers et al. 2003](#)):

- The total source spectrum consists of two additive components: one (the “burst component”) arising from nuclear burning and the other (“persistent emission”) arising from accretion.
- The burst component has the same spectral shape for all bursts from all sources.
- That shape is a blackbody, with variable temperature and normalisation.
- The persistent emission’s spectral shape does not change during a burst and is identical to its pre-burst shape.
- The persistent emission’s intensity does not change during a burst and is identical to its pre-burst level.

Making the last two of these assumptions allows the subtraction of the pre-burst emission of the neutron star from the burst spectra as background (e.g., [van Paradijs & Lewin 1986](#), [Lewin et al. 1993](#), [Kuulkers et al. 2003](#), [Galloway et al. 2008](#)). This approach (hereafter referred to as the “standard approach”) implicitly assumes that the accretion rate remains constant throughout a burst, but it is not obvious that this assumption is reasonable. For instance, when the flux reaches the Eddington limit one would naïvely expect accretion to cease entirely because the outward radiation force exceeds the gravitational force. [Lamb & Miller \(1995\)](#) argue that accretion ought to be shut off if the luminosity exceeds Eddington anywhere in the accretion flow, not necessarily just at the stellar surface. The effect of radiation halting accretion flow has apparently been observed for a number of the very brightest and most vigorous bursts (e.g., [in’t Zand & Weinberg 2010](#)), during which, for a few seconds, no X-ray flux (beyond instrumental background) is observed from the star—though this may also be due to the atmosphere expanding so far that it obscures the emitting regions of the accretion disc. On the other hand, increased luminosity during a burst might *enhance* the accretion rate, via radiation drag acting upon the accretion disc ([Walker & Meszaros 1989](#), [Miller & Lamb 1996](#), [Walker 1992](#), hereafter W92). At luminosities greater than $0.01L_{\text{Edd}}$, radiation forces have more of an effect on the accretion flow than general relativistic effects ([Miller & Lamb 1993](#)). Since it is not known whether the burst emission is isotropic or whether it is preferentially beamed in certain directions (e.g., [Boutloukos et al. 2010](#)), it is not obvious which of outwards pressure and radiation drag will dominate. It is therefore important to determine if varying accretion rates are detectable during a PRE burst, and to quantify any variation that is detected.

Some of the implicit assumptions of the standard analysis approach have been tested in previous studies: [van Paradijs & Lewin \(1986\)](#) pointed out that, if the total spectrum and persistent spectrum are both treated as blackbodies, subtracting the former from the latter will not leave a net burst spectrum that can be fit with a blackbody. This idea was followed up by [Kuulkers et al. \(2002\)](#) in a study of the high \dot{M} source GX 17+2. They did not find that accounting for this effect improved the spectral fits and concluded that the persistent emission does not originate from the same location as the burst emission on that neutron star. [Muno et al. \(2000\)](#) and [Strohmayer & Brown \(2002\)](#) allowed for the accretion to shut off entirely, by subtracting the instrumental background only, but

did not find that the spectral fits were improved by doing so. Recently [in't Zand et al. \(2013\)](#) studied a PRE burst from SAX J1808.4–3658 using combined *Chandra* and *RXTE* data and found that an observed excess of photons at both low and high energies can be well described by allowing a 20-fold increase of the pre-burst persistent emission. It also may be that the persistent emission is composed of separate contributions arising from different sites. These include a boundary layer at the inner edge of the accretion disk (e.g., [Kuulkers et al. 2002](#)), the inner regions of the accretion disk proper (e.g., [Christian & Swank 1997](#), [Cackett et al. 2010](#)), emission from the neutron star itself or its photosphere ([van Paradijs & Lewin 1986](#)), and Compton scattering in an accretion disk corona ([White & Holt 1982](#)). However, disentangling these contributions is likely to be difficult, because they are probably correlated and spectrally indistinct.

1.4.1 Spectral shape changes

Deviations of the burst component of the spectrum from a black-body spectrum could also be present. Such spectral changes have been theoretically predicted at both the high energy (e.g., [Pringle & Rees 1972](#), [London et al. 1984, 1986](#)), and low energy (e.g., [Madej 1991](#), [Madej et al. 2004](#)) ends of the X-ray spectrum. However, the literature is divided as to whether these are actually present in observations. Excess photons at high energy have been reported in bursts from 4U 2129+11 ([van Paradijs et al. 1990](#)) and GX 17+2 ([Kuulkers et al. 2002](#)). On the other hand, pure black-bodies have been found to give generally good results up to the present time (e.g., [Güver et al. 2012](#)) and some authors argue that they are consistent with black-bodies to extremely high confidence ([Boutloukos et al. 2010](#)). If the neutron star is spinning rapidly, then rotational broadening will affect the shape of the spectrum and can cause discrepancies of up to 10% in the inferred radius of the star ([Bauböck et al. 2015](#)). Additionally, material in the disc can become trapped at the corotation radius and is released in discrete episodes once enough has accumulated there ([D'Angelo & Spruit 2012](#)).

One might further divide the burst emission into contributions from a continuum, and discrete spectral features superimposed upon it such as emission lines and absorption edges. Continuum changes are likely to be present throughout all stages of a burst (e.g., [Suleimanov et al. 2011](#)). Changes in the spectral features are thought to be largely confined

to the Eddington-limited radius expansion period as these are thought to be due to ashes from nuclear burning being mixed into the expanding envelope (Weinberg et al. 2006), and have been detected in the so-called superexpansion bursts (in't Zand & Weinberg 2010). However, the non-detection of spectral features in a PRE burst from SAX J1808.4–3658 by in't Zand et al. (2013) suggests that such features may be too weak to be detected by currently available instruments: this source is the second-brightest PRE burster (behind only Cen X–4, which has been in quiescence since 1979, e.g., Chakrabarty et al. 2014) and the burst in question was observed by two X-ray observatories. Furthermore, Galloway et al. (2010b) had previously used *Chandra* spectra to search for spectral features in PRE bursts from 4U 1728–34, without success. Either the sought-after features are too weak to be detected in normal, frequent bursts, or they are not present in every burst.

Disentangling the burst component of the spectrum from the persistent component is difficult, as the various components are usually spectrally degenerate. Furthermore, observers typically do not have spatial resolution and therefore cannot attribute different spectral components to different sites in the neutron star system. This problem will be further complicated if, as expected, both the persistent and burst components can vary in spectral shape during a burst. Indeed, it has long been known that the spectral shape of the persistent emission, as well as its intensity, is a function of accretion rate (Hasinger & van der Klis 1989; see also Lyu et al. 2014). There are, however, suggestions that the accretion spectrum remains constant in shape on timescales longer than those of a burst, though its intensity may vary. During a burst it is, of course, much more difficult to detect changes in the accretion spectral shape. A sudden loss of high energy ($\gtrsim 30$ keV) photons has been detected from several sources, but this effect has been attributed to rapid cooling of the corona and is not related to accretion rate (Maccarone & Coppi 2003, Ji et al. 2013, Chen et al. 2013). Changes in the shape and intensity of the persistent emission during a very long (20,000 s) burst from 4U 1636–536 were reported by Keek et al. (2014), but variations on these timescales can happen on their own and the bursts I consider in this work are very much shorter.

1.5 Computer simulations

The first attempt at constructing numerical models of X-ray bursts was by [Joss \(1978\)](#). This model was a spherically symmetric, nonmagnetic and nonrelativistic atmosphere in which accretion occurred at the outer boundary. Helium burning was dealt with through a simple nuclear reaction network. Despite its simplicity, this model succeeded in reproducing the correct burst rise times (~ 0.2 s) and even modest radius expansion that lifted the photosphere to 30m above the burning layer.

Another simple model was the one-zone time-dependent model of [Paczynski \(1983a\)](#). This one allowed for nuclear burning but treated the atmosphere very simplistically, for example replacing the integrals of the atmosphere structure equations with their values at certain points. Nevertheless this model was useful because it required very little computing power, was easy to explain and gave sensible results.

[Taam \(1982\)](#) considered a scenario with accretion nuclear burning in a hydrogen-rich shell, and solved for a steady state solution in order to study the structure of the envelope. The nuclear network used was fairly sophisticated and included the $^{15}\text{O}(\alpha, \gamma)^{19}\text{Ne}$ reaction which removes a catalyst from the CNO cycle (see [Figure 1.4](#)) and has attracted considerable recent interest (e.g., [Fisker et al. 2007](#)). Another study around the same time by [Ayasli & Joss \(1982\)](#) incorporated an extensive nuclear network into the relativistic equations of stellar structure (see [Thorne 1977](#)). Their time-dependent models showed that the core temperature of the neutron star plays an important role in the structure of the atmosphere between bursts, and pointed out that the results of nuclear burning can influence the properties of subsequent bursts.

An early simulation that considered super-Eddington luminosities was that of [Kato \(1983\)](#). This model was also steady-state and spherically symmetrical, and treated the expanding atmosphere as a steady mass loss. Kato found that the total luminosity never exceeds the Eddington luminosity by more than a few percent, and that in the outer layers the product of luminosity and opacity is approximately constant. A similar quasi-static wind model was constructed by [Joss & Melia \(1987\)](#), which used more sophisticated boundary conditions and atmospheric opacity contributions, and obtained similar results.

[Paczynski \(1983b\)](#) considered helium burning at the base of a hydrogen-poor atmosphere and its effect on the structure of the atmosphere over time, using a modified stellar evolution code. This model took photon absorption and temperature-dependent electron scattering opacities into account, incorporated general relativity, used a very simple nuclear reaction network, and succeeded in producing radius expansion bursts on several occasions. It was found that only a small fraction of the total mass of the atmosphere expands; the innermost regions are hot enough to significantly reduce the electron scattering opacity. This allows super-Eddington flux to escape the inner regions, where it is efficiently converted to kinetic energy in the cooler and more opaque outer layers.

The work of London and collaborators ([London et al. 1984, 1986](#)) built upon the “classical” stellar atmosphere formulated by [Mihalas \(1978\)](#). Since the equilibration timescales are much shorter than the observed burst time scales, these authors argued that a quasi-equilibrium state exists for any given time and solved for it by iterating the equations of hydrostatic and radiative equilibrium. By varying the flux from the star, different stages of X-ray bursts could be investigated. Due to the assumption of an equilibrium state, only sub-Eddington fluxes could be considered in these models. These simulations generated X-ray spectra which included the effects of inelastic Compton scattering, and free-free and bound-free emission and absorption (see [Figure 1.8](#)). These spectra showed significant deviations from a blackbody. Other work by [Paczynski & Anderson \(1986\)](#) on static atmospheres investigated the effect of including general relativity and the temperature dependence of electron scattering and found that these things were essential factors in the structure of extended atmospheres, particularly when the flux approaches the Eddington limit.

Later computer models, such as those by [Madej \(1991\)](#), retain the assumption of equilibrium and increase the sophistication of the opacity effects in the atmosphere. Deviations from a blackbody spectrum in these works manifest as an excess of photons at the high energy end, consistent with the theoretical predictions of [Ebisuzaki et al. \(1984\)](#).

Considerable progress in simulating the burning layer have been made in recent years. For a rapidly spinning neutron star, [Spitkovsky et al. \(2002\)](#) have argued that Coriolis effects can cause the burning areas to take on the form of rotating Jupiter-like vortices. This may

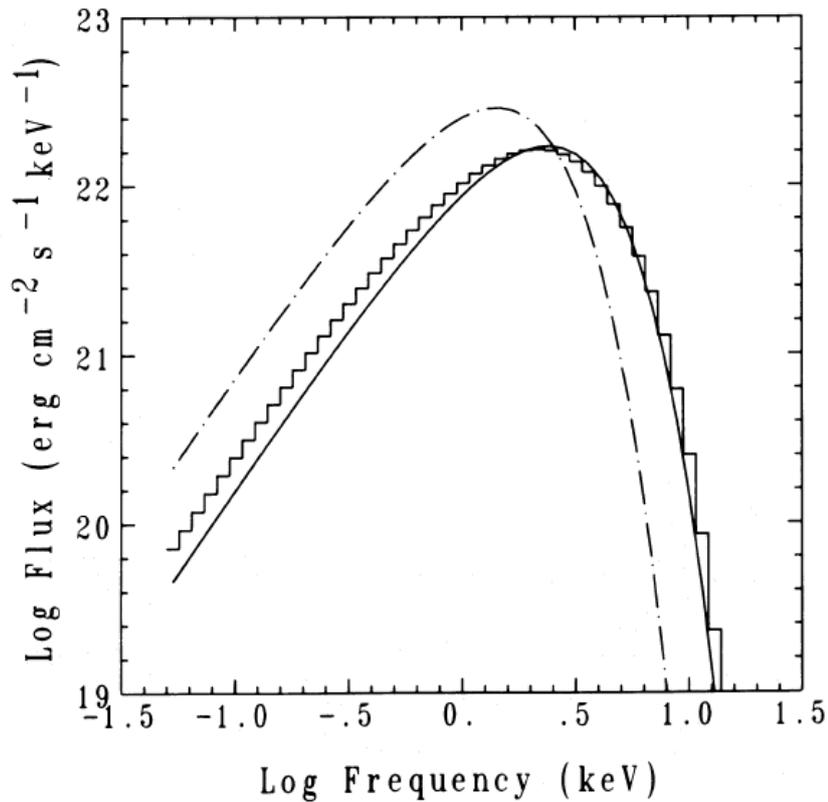


Figure 1.8 A simulated X-ray spectrum from one of the models of [London et al. \(1986\)](#). The stepped curve is the simulated spectrum, the dashed curve is a blackbody spectrum of the same effective temperature (0.5 keV), and the solid curve is the blackbody spectrum that best fits the simulated spectrum. The excess of high energy photons is evident.

be one explanation for QPOs. [Cavecchi et al. \(2013\)](#) performed hydrodynamic simulations of burning fronts and obtained rise times similar to those observed in X-ray bursts. They also found that flame spreading is faster for slower neutron star rotation and higher vertical thermal conductivity in the fuel layer.

The KEPLER code ([Weaver et al. 1978](#)) has allowed extensive studies of nuclear burning processes, and can reproduce type I bursts accurately as well as the period between bursts (e.g., [Heger, Cumming, Galloway & Woosley 2007](#)). This approach has revealed a new regime, where marginally stable burning produces quasi-periodic oscillations with long periods ([Heger, Cumming & Woosley 2007](#)). KEPLER simulations have produced a large catalogue of simulated bursts, which promises a more detailed comparison with real bursts.

1.6 Accretion disc structure

The interactions between a thermonuclear burst and the accretion flow itself have received less attention than the burning layers and the neutron star atmosphere. Let us discuss the structure of a geometrically thin disc. Consider portions of gas, with cylindrical coordinates (R, ϕ, z) . Since the disc is geometrically thin, all the particles lie near the $z = 0$ plane. It is usual to assume that the gas orbits on circular Keplerian paths, since for a given angular momentum a circular path has the least energy (e.g., [Pringle 1981](#)):

$$v_\phi = \sqrt{\frac{GM_*}{R}}. \quad (1.6.1)$$

Additionally the gas is allowed to move inwards or outwards with a radial velocity $v_R \ll v_\phi$ so that the orbits are still very nearly Keplerian and circular. Since the star is accreting from a companion I expect net inward flow (i.e., v_R negative) at most locations in the disc, and certainly near the neutron star.

The disc is characterised by its density, which is a function of its location in space, and can vary with time: $\rho(R, \phi, z, t)$. Similarly, v_R is also a function of position and time. It is typical to assume that the disc is geometrically thin and axisymmetric, which removes the dependence on ϕ, z . The radial velocity simply becomes $v_R(R, t)$ and I remove z from the density expression by introducing the *surface density*:

$$\Sigma(R, t) = \int_{-\infty}^{\infty} \rho(R, z, t) dz. \quad (1.6.2)$$

An annulus of disc gas located between R and $R + \Delta R$ has a total mass of

$$M(R, t) = 2\pi\Sigma(R, t)R\Delta R. \quad (1.6.3)$$

The rate of change of the mass is given by the flow from adjacent annuli:

$$\frac{\partial M}{\partial t} = v_R(R, t)M(R, t) - v_R(R + \Delta R, t)M(R + \Delta R, t), \quad (1.6.4)$$

which, if ΔR is small, can be written as

$$\frac{\partial M}{\partial t} \approx \Delta R \frac{\partial(Mv_R)}{\partial R}. \quad (1.6.5)$$

If I let $\Delta R \rightarrow 0$, and break $M(R, t)$ back up into R and Σ , I get the mass conservation equation

$$R \frac{\partial \Sigma}{\partial t} + \frac{\partial}{\partial R} (R \Sigma v_R) = 0. \quad (1.6.6)$$

The angular momentum conservation equation is of a similar form. Each annulus has angular momentum $M(R, t)R^2\Omega$, where Ω is its angular velocity. This quantity also flows between adjacent annuli, and it is also necessary to include a contribution arising from viscous torques exerted between them. Note that

$$\Omega(R) = \left(\frac{GM_*}{R^3} \right)^{1/2} \quad (1.6.7)$$

is a function of radius only, since it is the angular velocity of Keplerian orbits at that radius, and this does not change with time. I have

$$\begin{aligned} \frac{\partial}{\partial t} [M(R, t)R^2\Omega] &= v_R(R, t)M(R, t)R^2\Omega(R) \\ &- v_R(R + \Delta R, t)M(R + \Delta R, t)(R + \Delta R)^2\Omega(R + \Delta R) \\ &+ \frac{\partial G}{\partial R}\Delta R, \end{aligned} \quad (1.6.8)$$

where $G(R, t)$ is the viscous torque exerted by the outer annulus on the inner. The torque exerted by the inner on the outer is, obviously, $-G(R, t)$. It may be time-dependent because it may involve quantities such as Σ , which are themselves time-dependent. It can therefore be seen that viscous processes tend to spin down inner annuli and spin up outer ones.

Viscous torque is thought to arise from turbulent motions within the disc, but the details are still not well understood. Turbulence can be thought of as portions of gas being exchanged between adjacent annuli in equal quantities. These local motions are distinct from the global radial motion of gas described by v_R . Although exchanging two equal quantities of gas between annuli cause no net flow of mass, there is a net flow of angular

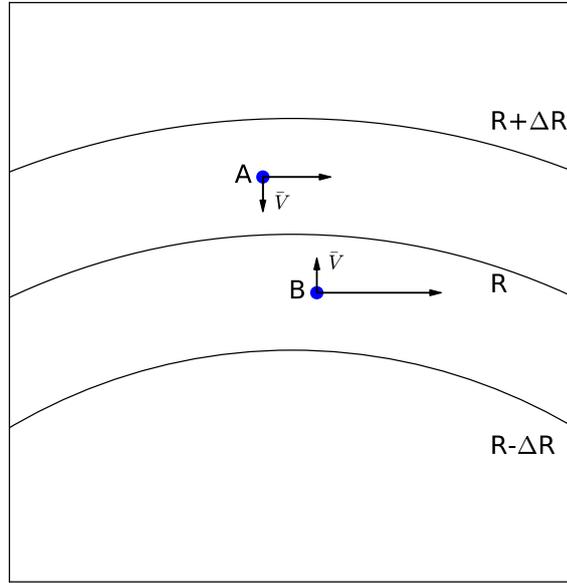


Figure 1.9 Two identical portions of gas (A and B) being exchanged, with radial velocity \bar{V} , between annular rings of an accretion disc. The inner portion, B, has greater angular velocity than A. Therefore there is a net flow of angular momentum outward, resulting in torques between annuli, even though there is no net flow of mass.

momentum. Viscosity is the cause of accretion— without it, the gas would continue on Keplerian orbits forever.

Suppose two equal packets of gas near radius R are exchanged, each moving a distance ΔR with a speed \bar{V} . That is, one packet (which I label A) moves from $R + \Delta R/2$ to $R - \Delta R/2$, while the second one (B) moves the other way (see Figure 1.9). From the point of view of an observer located at radius R and corotating with the disc at that point, with angular velocity $\Omega(R)$, packet B moves with velocity $(R - \Delta R/2)\Omega(R - \Delta R/2) + \Omega(R)\Delta R/2$, and a similar relation holds for packet A .

Since these processes are occurring all round the circumference of the annulus, I can write down the flux of angular momentum outward through a unit arc length:

$$\frac{dL(R)}{dt} = \Sigma(R, t)\bar{V}(R - \Delta R/2), \quad (1.6.9)$$

and similar for the gas moving in. The torques exerted by outer annuli on inner ones are given by the net angular momentum flux over the entire boundary, which is, to first order

$$G(R) = 2\pi R^3 \Sigma \nu \frac{d\Omega}{dR}, \quad (1.6.10)$$

where the kinematic viscosity $\nu = \bar{V} \Delta R$.

It follows that

$$R \frac{\partial}{\partial t} (\Sigma R^2 \Omega) + \frac{\partial}{\partial R} (R^3 \Sigma \nu_R \Omega) = \frac{1}{2\pi} \frac{\partial G}{\partial R}, \quad (1.6.11)$$

and I can combine this with equations 1.6.6, 1.6.7, and 1.6.10 to get the time evolution of surface density (see Pringle 1981, Frank et al. 2002 for details):

$$\frac{\partial \Sigma}{\partial t} = \frac{3}{R} \frac{\partial}{\partial R} \left(R^{1/2} \frac{\partial}{\partial R} \left[\nu \Sigma R^{1/2} \right] \right). \quad (1.6.12)$$

Accretion from a companion star may be accounted for by adding a source term $S(R, t)$ to this equation (Bath & Pringle 1982):

$$\frac{\partial \Sigma}{\partial t} = \frac{3}{R} \frac{\partial}{\partial R} \left(R^{1/2} \frac{\partial}{\partial R} \left[\nu \Sigma R^{1/2} \right] \right) + S(R, t), \quad (1.6.13)$$

but I will not incorporate such a term in this work. In Chapter 5 I compare three-dimensional simulations of a non-irradiated accretion disc to a one-dimensional code that integrates Equation 1.6.12, to verify the correctness of the simulations.

The magnitude of the kinematic viscosity ν is difficult to estimate. It is likely that the longest turbulent lengthscales will be shorter than the physical height of the disc (expressed as *scale height* $H(R)$), and it is unlikely that turbulence will exceed the local sound speed c_s ; otherwise, shocks would probably convert the supersonic turbulent motions into heat before very long. Thus, we can encapsulate all our ignorance of turbulent processes in the disc into a single parameter α (Shakura & Sunyaev 1973):

$$\nu = \alpha c_s H. \quad (1.6.14)$$

On dimensional grounds, α is unitless. We expect it to be less than unity, but we are ignorant as to its true value. It may not even be constant throughout the structure of the

disc. Observational evidence suggests it is typically in the range $0.1 \lesssim \alpha \lesssim 0.4$ (King et al. 2007). A further source of viscosity is *bulk viscosity*, arising from internal friction when a fluid is compressed or expanded. In the accretion discs I consider, rapid compression or expansion are unlikely so this source of viscosity can be neglected.

1.6.1 Vertical disc structure

In a geometrically thin disc the vertical and radial structures are decoupled. Though the self-gravity of the disc can be neglected, it must still be in hydrostatic equilibrium because the gravitational pull from the star has a vertical component, so

$$\frac{1}{\rho} \frac{\partial P}{\partial z} = \frac{\partial}{\partial z} \left(\frac{GM_* z}{\sqrt{R^2 + z^2}} \right). \quad (1.6.15)$$

We take the disc to be isothermal (so that $P \propto \rho$) and, because it is geometrically thin, $z \ll R$. Therefore I can write

$$c_s^2 \frac{\partial [\log_e \rho]}{\partial z} = -\frac{GM_* z}{R^3}, \quad (1.6.16)$$

where $P = c_s^2 \rho$. Solving this for density gives

$$\rho(R, z) = \exp \left[\frac{-GM_* z^2}{2c_s^2 R^3} \right]. \quad (1.6.17)$$

The *scale height* H of the disc is defined as

$$H^2 = c_s^2 \left(\frac{R^3}{GM_*} \right) = \left(\frac{c_s}{\Omega} \right)^2. \quad (1.6.18)$$

Thus the vertical structure of the disc is given by:

$$\rho(R, z) = \rho_0(R) \exp \left[\frac{-z^2}{2H(R)^2} \right], \quad (1.6.19)$$

where, $\rho_0(R)$ is the density at the midplane. But since

$$\Sigma(R) = \int_{-\infty}^{\infty} \rho(R, z') dz', \quad (1.6.20)$$

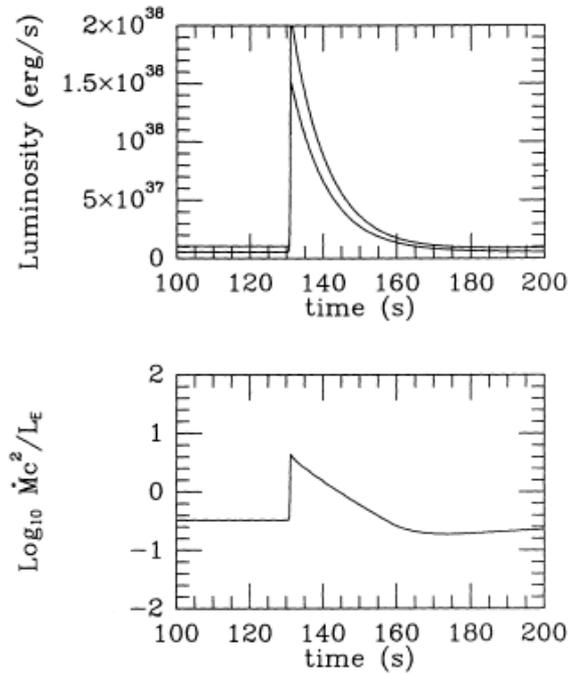


Figure 1.10 The result of one simulation from Walker (1992). The top panel shows the luminosity of the star; the bottom panel shows the mass accretion rate. There is an almost instantaneous increase in accretion rate of over one order of magnitude at $t \approx 30$ s, corresponding to $f_a = 12$ (see Walker’s Table 1). A slight decline in accretion rate to slightly below the pre-burst level is also evident at about 170 s. It is due to the disc being depleted of material by the preceding episode of enhanced accretion.

it is possible to eliminate $\rho_0(R)$:

$$\frac{\Sigma(R)}{\sqrt{2\pi}H(R)} \exp\left[\frac{-z^2}{2H(R)^2}\right]. \quad (1.6.21)$$

1.6.2 Previous theoretical analyses

W92 performed one-dimensional simulations of geometrically thin, axisymmetric irradiated accretion discs around a neutron star. Their models include the effects of viscosity and general relativity. W92’s results predict that radiation torque from bursts can enhance the accretion rate by up to two orders of magnitude; see Figure 1.10.

The evolution of disc surface density Σ in W92 is a general-relativistic analogue of 1.6.12:

$$\frac{\partial \Sigma}{\partial t} = \frac{1}{u^t R} \frac{\partial}{\partial R} \left\{ \frac{R}{u_{\phi,R}} \left[\frac{1}{R} \frac{\partial}{\partial R} (\gamma e^{\psi - \lambda} r W) + \int dz k_{\phi} \nabla \cdot \mathbf{S} \right] \right\}. \quad (1.6.22)$$

In this expression u is the four-velocity, W is the shear stress integrated over the vertical extent of the disc, and γ is the Lorentz factor. The vertical integral describes angular momentum loss in terms of the photon flux four-vector \mathbf{S} and the transverse component of the momentum four-vector \mathbf{k} .

In these simulations the vertical structure of the disc was not considered; instead, W92 assumed that radiation forces would be efficiently communicated to the interior so that the vertical dimension can be neglected. The disc was illuminated by a simulated burst consisting of an exponential rise in intensity over 0.1 s followed by an exponential decay over 10 s. W92 performed these simulations on stars with different pre-burst accretion rates, spin periods, and stellar radii and found that significant enhancements in accretion rate are possible in all cases. The effect is smaller for more quickly rotating stars, and there is also a slight increase with decreasing stellar radius.

1.7 Outline of this thesis

In this thesis I challenge the assumption that the intensity of persistent emission remains constant throughout a burst. I argue that radiation forces from a thermonuclear burst irradiating the disc could influence the accretion rate, and this could potentially be detected as a change in the overall burst spectrum. Such radiation forces could either induce a decrease in accretion, due to outwards radiation pressure inhibiting the inflow of material, or increase it via the torque exerted on the accretion disc moving through a radiation field—the so-called *Poynting-Robertson effect* (Poynting 1904, Robertson 1937). Estimates by Ballantyne & Everett (2005) suggest that the timescale of radiation drag effects on the disc are of the order of a few seconds, so these should be detectable in the data.

Computer simulations of the burning layer and atmosphere have reached a high level of sophistication. Many phenomena observed from bursters have been reproduced accurately in simulations, and great physical insight has been gained. The behaviour of the accretion disc in response to a burst has not yet been modelled to such a detailed extent. Walker (1992) studied radiation drag on the accretion disc in a one-dimensional, relativistic simulation, finding that large enhancements are possible, but there have as yet been no three-dimensional computational studies. This thesis aims, in part, to fill this gap.

Chapter 2

Methods and Data Analysis

For the initial attempt to determine whether the intensity of persistent emission can change during a burst, I have selected 332 PRE bursts from 40 individual neutron stars. I selected this subset of bursts because PRE bursts are the brightest from any given source and therefore yield the best signal to noise, and because PRE bursts may give the best chance of seeing all the possible accretion increase or decrease behaviours.

2.1 Data collection and reduction

2.1.1 Introduction to X-ray spectral analysis

Most spectroscopy of X-ray bursts is done with proportional counters or scintillation counters, which give photon count rates in a range of pulse-height-analyzer (PHA) channels (Lewin et al. 1993). A proportional counter consists of a gas mixture enclosed within an airtight conducting container. The gas is typically about 90% of a noble gas and 10% polyatomic gas, maintained at about one atmosphere of pressure. One side of the container is a thin sheet of plastic or light metal, called the *window*, through which the X-rays enter. The vessel also contains an anode wire held at about 2,000 V. An X-ray entering the detector strikes one of the noble gas atoms and produces a photoelectron with energy similar to that of the incident photon. This loses its energy by ionising further atoms and producing more electrons which are drawn to the anode wire, gaining enough energy as they are accelerated to produce even more electrons. In this way a single X-ray photon

can be amplified to a detectable level, and in a manner that its original energy can be measured. Most detectors in the 0.2-20 keV range are proportional counters (Adams 1980). In front of the window is an arrangement of parallel metal tubes called the *collimator*. Its function is to block X-rays arriving at the detector from angles far away from the area of sky under observation. It can also be used to push against the window, supporting it against the pressure difference between the gas in the proportional counter and the vacuum outside.

The relationship between incident photon spectrum and photon counts in individual PHA channels is not trivial. Detectors have energy-dependent sensitivity, and a detected photon may be recorded in any one of several PHA energy bins. These processes can be described by an *ancillary response matrix*, which describes the likelihood that a photon with a given energy will be detected at all, and a *redistribution matrix*, which describes how photons that are detected are assigned to the PHA channels.

The *ancillary response* matrix quantifies the sensitivity of the detector as a function of photon energy:

$$S(E) = A \times P(E), \quad (2.1.1)$$

where A is the collecting area of the detector in units of cm^2 and $P(E)$ is the probability that a photon of energy E entering the instrument will be detected by it. Formally,

$$P(E) : \mathbb{R}_+ \rightarrow [0, 1]. \quad (2.1.2)$$

The *redistribution matrix* is given by

$$R(E) : \mathbb{R}_+ \rightarrow \mathbb{Z}^n, \quad (2.1.3)$$

where n is the number of PHA channels in the instrument. That is, to each photon energy is assigned a vector whose elements are the probabilities that a photon of that energy will be assigned to the corresponding PHA channels. In practice these functions are not represented as continuous functions of photon energy, but are discretised into small energy bins.

It is possible to combine the ancillary response and redistribution matrices into one *response matrix*, but doing so is not required. Inverting this matrix would give the incident photon spectrum in terms of the photon counts in PHA channels, but performing this operation is not a numerically stable process (Blissett & Cruise 1979). It is therefore not possible in general to infer the source spectrum from detected photon counts.

Instead, it is typical to hypothesise a source spectrum, with a number of variable parameters, and pass it through the response matrix to obtain a predicted count rate spectrum. This can be compared to the observed counts and the suitability of the hypothesised spectrum evaluated by some goodness-of-fit estimator, typically χ^2 . The estimator may also be modified to give less weight to channels with low counts (e.g., Churazov et al. 1996).

2.1.2 Burst spectral analysis with RXTE

I used observational data from the *Rossi X-Ray Timing Explorer* (*RXTE*), publicly available through the High-Energy Astrophysics Science Archive Research Centre (HEASARC),¹ dating from shortly after the satellite's launch on December 30, 1995 to the end of the *RXTE* mission on January 3, 2012. *RXTE* carries three instruments for detecting X-rays. The All-Sky Monitoring camera (ASM) consists of three scanning shadow cameras sensitive to photons with energies between 1.5 and 12 keV with a field of view of approximately two degrees, and performed 90s step-stare observations of most of the sky every 96 minutes (Levine et al. 1996). The Proportional Counter Array (PCA) consists of five proportional counter units (PCUs), with effective areas of 1,600 cm² each, sensitive to photons with energies between 2 and 60 keV and has a field of view of approximately one degree (Jahoda et al. 2006). *RXTE* also carries the High Energy X-ray Timing Experiment (HEXTE), a collection of scintillation counters with a one degree field of view (Rothschild et al. 1998), but I do not use data from HEXTE in this work.

The sample of radius expansion bursts is based on the burst catalogue of G08, with the addition of 118 type I bursts detected after the publication of that paper. This sample forms part of the Multi-INstrument Burst Archive (MINBAR)². The G08 sample distinguishes between PRE bursts and non-PRE bursts by defining a PRE burst as one that i) reaches a

¹See <http://heasarc.gsfc.nasa.gov>

² see <http://burst.sci.monash.edu/minbar>

local maximum blackbody normalisation K_{bb} at or near the moment of peak flux, ii) has declining values of K_{bb} after this time, and iii) attains its lowest blackbody temperature at the maximum K_{bb} (see Figure 2.3. A relatively small number of bursts (58) meet some, but not all, of these criteria. These are described as partial radius expansion bursts, and I exclude them from this work.

Unless otherwise stated, the data analysis procedures are as in G08. The catalogue contains time-resolved spectra in the range 2–60 keV covering the burst duration, for all bursts for which the suitable datamodes were available. These spectra were extracted on intervals, beginning at a length of 0.25 s during the burst rise and peak. To maintain the same signal-to-noise level during periods of declining flux, the bin size was gradually increased.

When an *RXTE* PCU detects a photon, it becomes inactive for a short ($\approx 10 \mu\text{s}$) interval. This is called “deadtime”, and it causes a small fraction of the incident photons to not be detected. For an incident photon count rate of $400 \text{ count s}^{-1} \text{ PCU}^{-1}$, deadtime causes a loss of about 3% of the photons. I calculated an effective exposure for each measured spectrum, that accounts the fraction of photon detection events lost during deadtime, as recommended by the *RXTE* instrument team³. Dead time contains contributions from both coincidence and particle events as well as photons from the source and background.

I re-fit the spectra over the energy range 2.5–20 keV using the revised PCA response matrices, v11.7⁴ and adopted the recommended systematic error of 0.5%. The fitting was undertaken using XSPEC version 12 (Arnaud 1996, Dorman & Arnaud 2001). In order to accommodate spectral bins with low count rates, I adopted Churazov weighting (Churazov et al. 1996), which estimates the statistical weight for a given channel by averaging photon counts in neighbouring channels.

I modelled the effects of interstellar absorption, using a multiplicative model component (`wabs` in XSPEC), with the column density N_{H} frozen. The N_{H} values are drawn from the literature, preferentially from studies of neutron stars using instruments sensitive at lower X-ray energies than *RXTE*. The values adopted are listed in Appendix A. In the original analysis carried out by G08, the neutral absorption was determined separately for each burst, from the mean value obtained for spectral fits carried out with the N_{H} value free to

³see http://heasarc.gsfc.nasa.gov/docs/xte/recipes/pca_deadtime.html

⁴see <http://www.universe.nasa.gov/xrays/programs/rxte/pca/doc/rmf/pcarmf-11.7>

vary with time. This has a negligible effect on the burst flux, but can introduce spurious burst-to-burst variations in the blackbody normalisation.

2.2 Method

The revised analysis is to fit the burst spectra with a two-component model consisting of a black-body, representing the burst emission, and a model of the pre-burst persistent emission, representing the emission due to accretion of material onto the neutron star, with a prefactor f_a left free in the fits. As I cannot distinguish contributions to the persistent emission arising from different locations in the neutron star system, I simply assume the persistent emission is indivisible and results entirely from accretion onto the neutron star. I include non-PRE bursts later, in Chapter 4.

It is also possible that the burst emission is non-Planckian, but fitting a non-black-body spectrum requires theoretical models that can describe the data. The most recent model atmospheres of bursting neutron stars are those of [Suleimanov et al. \(2011\)](#), but even these have been only partially successful ([Zamfir et al. 2012](#)), and are not valid for radius expansion burst spectra because they do not allow for the expansion of the atmosphere. In the absence of models that are demonstrably better than black-bodies, and no consensus that black-body fits really are unsuitable, I keep the assumption of thermal burst emission for the majority of the analysis. I discuss this issue further in §4.3. Similarly, I do not draw a distinction between deviations from a blackbody due to continuum or features in this work, since spectral features are either not present in every burst or are too weak to be detected.

In order to demonstrate the new approach, I selected a burst from the well-studied PRE burster 4U 1636–536 ([Swank et al. 1976](#)). This source is an attractive candidate for several reasons. It is a prolific burster, with 75 PRE bursts observed by *RXTE*. At a distance of approximately 6 kpc (G08) it is relatively nearby. Its average peak burst flux of $(69 \pm 6) \times 10^{-9}$ erg cm⁻² s⁻¹ places it among the brightest sources and ensures good signal-to-noise. The hydrogen column density of 0.25×10^{22} cm⁻² ([Asai et al. 2000](#)) is low compared to most other sources in the catalogue, minimising the absorption corrections that have to be performed on the spectra. There are no other known LMXBs in the same

field of view as 4U 1636–536, so confusion with the persistent emission of other sources is not an issue (see §2.2.3 for further discussion of this problem). Finally, 4U 1636–536 accretes mixed H/He (Galloway et al. 2006) and is therefore representative of the majority of neutron stars in the sample (G08).

The burst I chose was detected on Jun 15, 2001 (burst ID 34 in G08). The data for this burst consists of 122 spectra taken by *RXTE* with an integration time of typically 0.25s. Of these spectra, 37 are after the beginning of the burst, which I take to be the earliest time that the flux exceeds 25% of the maximum burst flux. The high time-resolution data continued up to 176.75s after the burst start, but the total flux had declined to the pre-burst level by 24.75s. Two of *RXTE*'s five PCUs, numbers 3 and 4, were active for this burst.

2.2.1 Characterising the persistent emission

I adopted the integrated X-ray flux for a 16-second interval prior to the start of each burst as the persistent emission. This spectrum also includes a time-dependent contribution from instrumental (non-source) background; to estimate this contribution I used the full-mission, “bright” source ($> 40 \text{ count s}^{-1}$) models released 2006 August 6 with the `pcabackest` tool in the FTOOLS package (Blackburn 1995). Subsequent model fits to each persistent (and burst) spectrum used the corresponding model spectrum estimated for that burst as background. I then adopted a model for the persistent emission consisting of a blackbody plus a power law, both corrected for interstellar absorption. The fits were performed with XSPEC (Arnaud 1996) using `wabs*(bbodyrad+powerlaw)` as the model. The hydrogen column density was kept fixed at $0.25 \times 10^{22} \text{ cm}^{-2}$, the value determined in Asai et al. (2000).

For the persistent emission model I obtained a blackbody temperature of $(1.9_{-0.1}^{+0.1}) \text{ keV}$ and a normalisation of $(6.0_{-1.0}^{+1.2}) \text{ km}^2 / (10 \text{ kpc})^2$. The powerlaw component had index $(+3.0_{-0.2}^{+0.3})$ and normalisation $(2.2_{-0.5}^{+0.8}) \text{ keV}^{-1} \text{ cm}^{-2} \text{ s}^{-1}$ at 1 keV. Errors given are the 1σ confidence level. The reduced χ^2 for this fit was 0.699 for 21 degrees of freedom, indicating that the model adequately describes the data ($p = 0.84$). Figure 2.1 shows the fit to the data and the residuals.

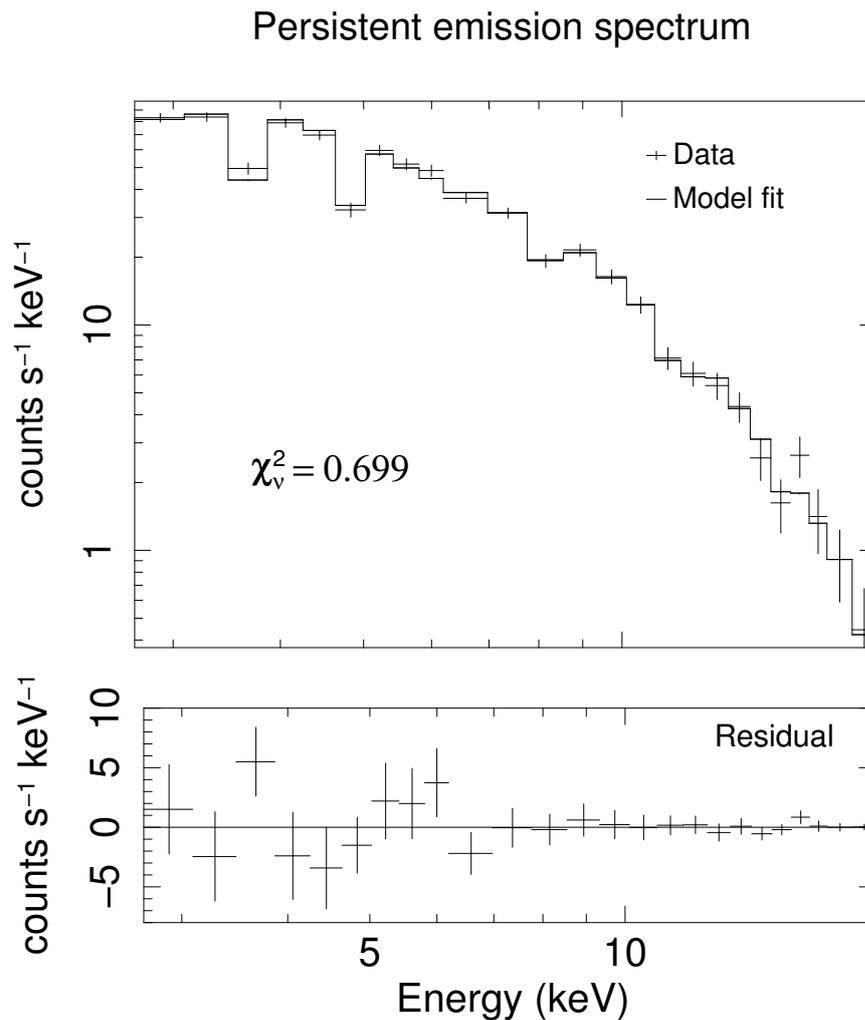


Figure 2.1 Measured persistent spectrum for the Jun 15, 2001 burst from 4U 1636–536 and the fitted spectral model (top panel), and residuals to the data (bottom panel). The reduced χ^2 is 0.699, for 21 degrees of freedom, indicating a good fit to the data.

2.2.2 Modelling the burst spectrum

I initially fit the net burst spectra with the standard approach: a blackbody spectrum, corrected for interstellar absorption using the same hydrogen column density as above, i.e. `wabs*bbbodyrad` in XSPEC. I subtracted the measured pre-burst emission, which includes a component that does not arise from the source (the instrumental background), and fit the resulting burst spectrum. This is the same implementation of the standard approach as in G08 and earlier works. These initial fits provide standard approach fits to which I can compare the new approach, as well as sensible spectral parameters to seed the variable persistent emission fits.

Now I introduce the dimensionless quantity f_a , the persistent emission multiplicative factor, as a third variable parameter. The burst spectra are fitted again, replacing the recorded pre-burst emission subtraction with just the instrumental background, with the model

$$S(E) = A(E) \times B(E; T_{\text{bb}}, K_{\text{bb}}) + f_a \times P(E) + b(E)_{\text{inst}}, \quad (2.2.1)$$

where $S(E)$ is the fitted spectrum as a function of energy E , A is the absorption correction, B is the blackbody function with temperature T_{bb} and normalisation K_{bb} , P is the persistent model described in §2.2.1, and b_{inst} is the instrumental background. Note that the persistent model already contains an absorption factor. The parameter f_a is allowed to vary between -100 and 100. Because of the 0.25 s integration time of the burst spectra, the measured peak f_a values may be slight underestimates; the instantaneous maximum in any time interval must be greater than the average over that interval.

Figure 2.2 shows the time evolution of burst bolometric flux (top panel), the f_a factor (middle panel), and the reduced χ^2 for both treatments (lower panel) for the selected burst. I found that the f_a increases to many times the pre-burst level, peaking at 17.8 ± 4.7 at 0.25 seconds after the burst start. The errors are 1σ significance limits determined by XSPEC. The burst bolometric flux is calculated from the blackbody parameters

$$\begin{aligned} F &= \sigma T_{\text{bb}}^4 (R/d)^2 \\ &= 1.076 \times 10^{-11} \left(\frac{kT_{\text{bb}}}{1 \text{ keV}} \right)^4 K_{\text{bb}} \text{ erg cm}^{-2} \text{ s}^{-1}, \end{aligned} \quad (2.2.2)$$

where T_{bb} is the blackbody temperature, R is the effective radius of the emitting surface, d is the distance to the source, and K_{bb} is the normalisation of the blackbody. The calculated flux does not include the contribution to the total flux due to the scaled persistent emission. The burst component flux for the variable persistent flux fits is therefore consistently lower than the standard approach fits for the peak of the burst, where $f_a > 1$. I also find that, after the beginning of the burst, the χ^2_{ν} for the variable persistent flux fits is consistently lower than for the standard approach fit, with means of 1.26 ± 0.59 and 1.50 ± 0.60 respectively. A Kolmogorov-Smirnov test shows that the two sets of χ^2_{ν} values have a 4.0% probability of being consistent with one another. Figure 2.4 shows the flux-temperature curve for this burst.

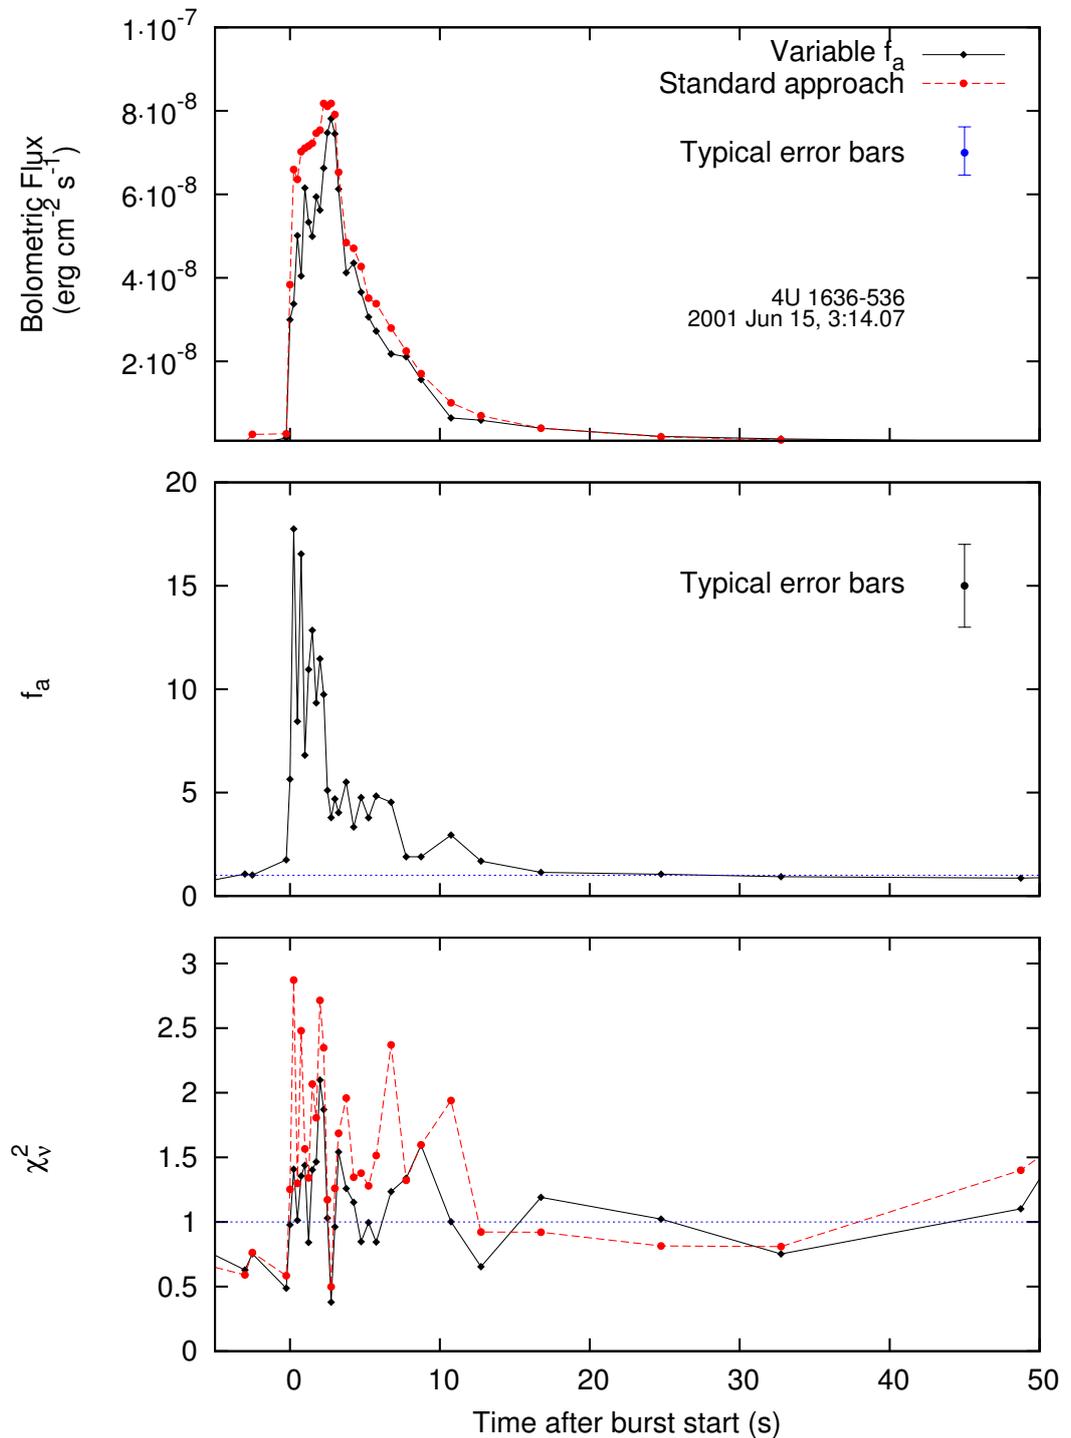


Figure 2.2 Comparison of fitting a variable persistent emission factor, f_a , to the standard approach fits for a burst from 4U 1636–536, where f_a is a scaling factor of the pre-burst emission. The variable fit yields consistently lower burst fluxes (top panel), though these are within the error bars. The contribution to the flux from the variable persistent emission increases to 17.75 times the pre-burst level (middle panel) during the rise of the burst. The f_a factor remains elevated for a considerable time after the total flux begins to decrease. Allowing the persistent emission to vary improves the reduced χ^2 (lower panel). See Figure 2.4 for the changes in normalisation and temperature for this burst.

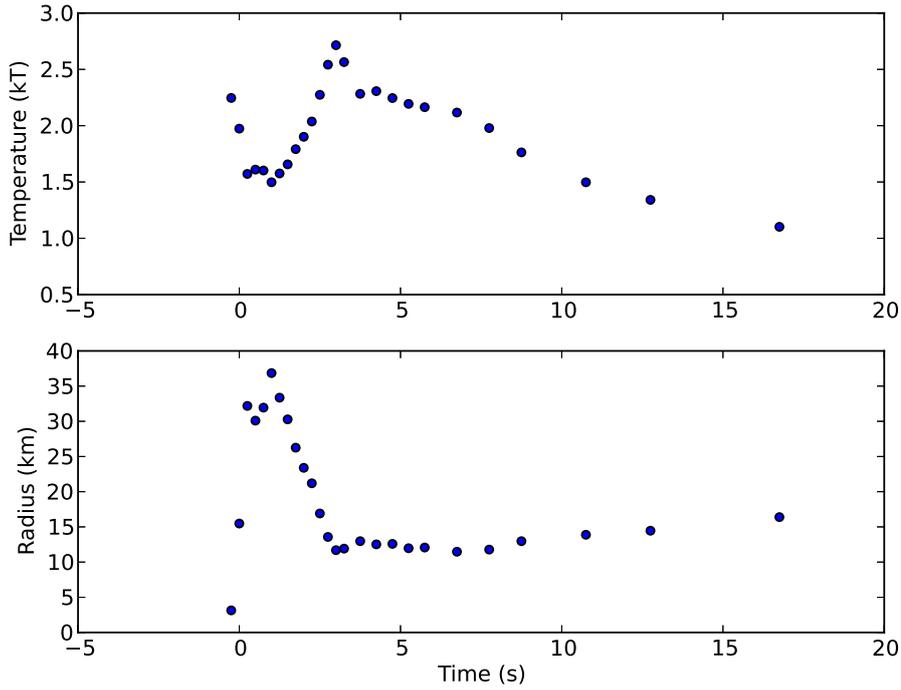


Figure 2.3 Blackbody temperature and radius for the "prototype" burst from 4U 1636–536, showing the major features of radius expansion (§2.1). This burst shows low temperature at the moment of maximum radius, indicative of the photosphere being expanded and cooled. Later the photosphere high temperature for low radius, followed by declining temperature at approximately constant radius. Such behaviour for temperature and radius is characteristic of radius expansion. For clarity I have not plotted error bars.

The bolometric flux of the pre-burst persistent emission lower than that of the brightest burst spectrum, by a factor of 35, so it is possible that counting statistics in the burst spectrum could induce a spurious response in f_a . To investigate this possibility I examined the spectrum with the highest f_a , which was 17.8 ± 4.7 measured 0.25s after the burst start. I took the blackbody parameters for that spectrum, $kT = 1.572$ keV and normalisation $1,034 \text{ km}^2 / (10 \text{ kpc})^2$, and generated 1,000 simulated spectra consisting of an absorbed blackbody with those parameters plus the persistent model times unity. The simulated spectra also incorporated counting statistics typical of the detector. These parameters represent a burst spectrum for which f_a does not change and for which any measured deviation from $f_a = 1$ must be due to noise or the fitting procedure.

I fit each of the simulated spectra with the new variable persistent flux model. The mean parameters for these were $kT = 1.569 \pm 0.011$ keV, $K_{\text{bb}} = 1,061 \pm 49 \text{ km}^2 / (10\text{kpc})^2$ and $f_a = 0.36 \pm 1.18$. The mean parameters are therefore all within 1σ of those that seeded the simulated fits. Adopting the standard deviation of the simulation values as the error on

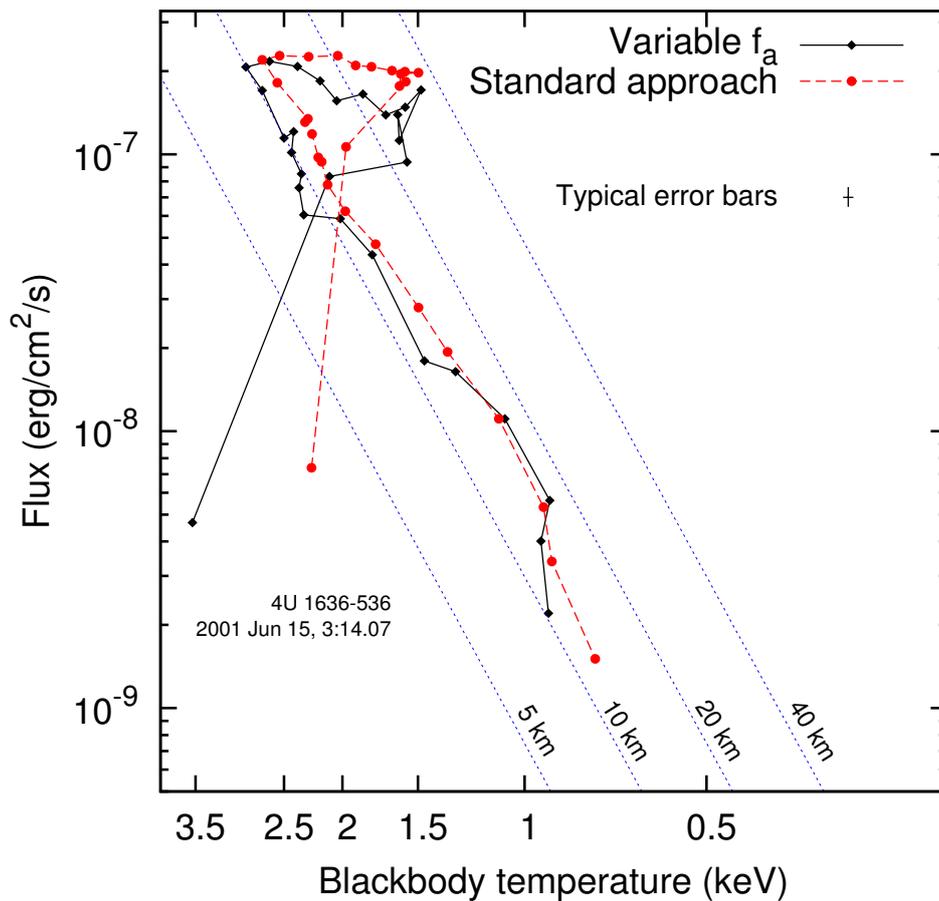


Figure 2.4 Flux-temperature diagram for the 2001 Jun 15 burst from 4U 1636–536, as measured by both the standard fit and the variable persistent flux fit. The fluxes shown are bolometric fluxes calculated from the blackbody parameters, assuming a distance to the source of 6 kpc (G08). The variable fit is consistently slightly lower in flux, though within the uncertainties. Also plotted are lines of constant radius for 5 km, 10 km, 20 km, and 40 km. These radii do not take into account spectral corrections (e.g., [Suleimanov et al. 2011](#)). The radius of the photosphere at maximum expansion is the same using both methods, but the variable persistent flux fit gives a smaller, hotter photosphere at the touchdown point and beginning of the cooling tail. Typical error bars are given for points near the time of peak flux; at lower fluxes they are much larger.

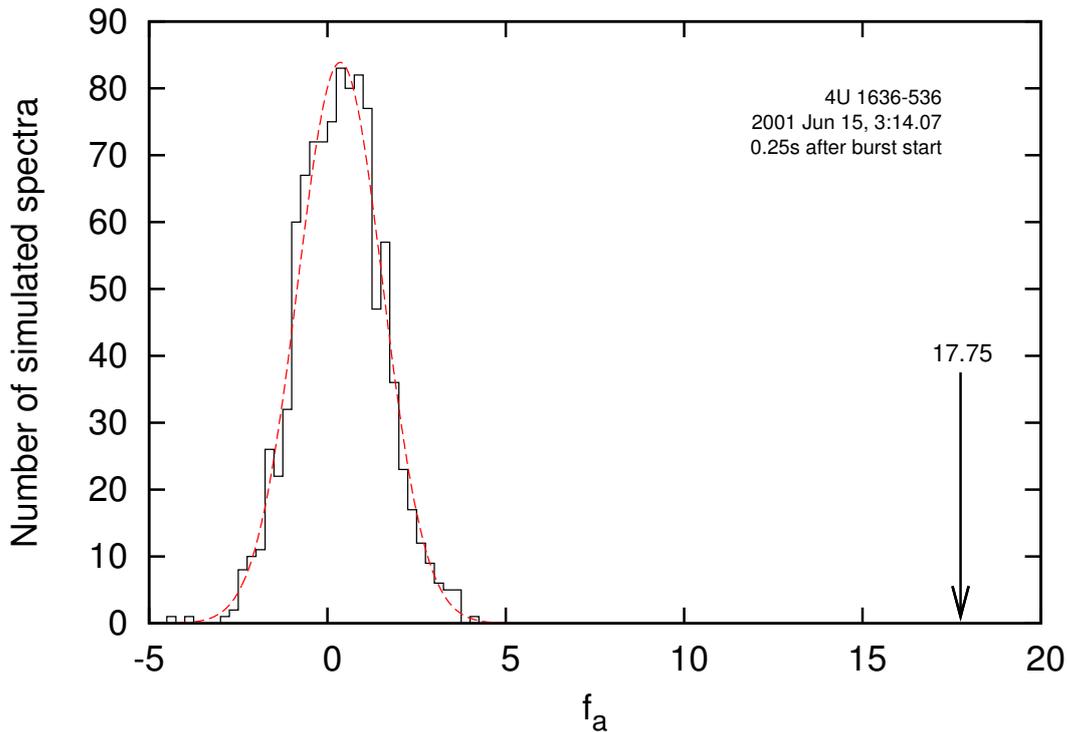


Figure 2.5 Histogram of the measured persistent emission factor f_a for 1,000 simulated spectra for the burst from 4U 1636–536. These spectra consist of a blackbody with parameters taken from the real spectrum with the highest measured f_a , plus the persistent model times unity. The scatter arises from the (simulated) counting statistics. The observed spectrum is 14.6 standard deviations from the predicted f_a assuming the persistent emission does not change. This shows that high measurements of f_a cannot be attributed to counting statistics.

the fitted f_a I found that the measured f_a for the real burst spectrum was 14.6 standard deviations higher than the simulated mean. I therefore conclude that the high f_a values do not arise from counting statistics alone. Figure 2.5 shows a histogram of the measured f_a for the simulated spectra compared with the value measured from the real burst spectrum. Thus, for one burst I have found evidence for a better spectral fits for a model which can be described by varying the level of the persistent emission. This change is statistically significant for at least some of the spectra; 20 of the 37 of the spectra after the beginning of the burst had f_a above unity with more than 4σ significance (using the above procedure, see also Figure 2.5). The associated improvement in χ^2 is of greater than 3σ significance using f-tests. I apply the new method to all the PRE bursts in the burst catalogue in the next section.

2.2.3 Burst selection

I restrict the sample of PRE bursts to exclude events that are unsuitable for analysis. Three bursts were excluded for which the Standard-2 data was missing, preventing estimation of the instrumental background via `pcabackest`.

Some burst sources, such as 4U 1728-34, AQL X-1, and GRS 1747–312, lie in crowded fields containing other LMXBs (see Figure 1.1). If the other source(s) were active at the time of observation, then their persistent emission could be confused with the source under observation and it would not be possible to scale only the persistent flux from the burst source. These sources need to be excluded from consideration.

I used All-Sky Monitor (ASM) data to quantify this effect⁵. For each burst I found every source within the PCA field of view and took the ASM counts at the time record nearest the burst start time. Since the PCA response drops off approximately linearly with distance from the field of view’s center, I multiplied the ASM counts by $1 - s$, where s is the angular distance of the source from the center of the field of view, in degrees. I compared all the other sources to the burst source, and if the total contribution is more than 10% of the pre-burst flux of the burst source I excluded this burst from consideration. I found source confusion in 36 bursts. The majority of these are from 4U 1728-34, which lies only 0.56° in the sky away from the Rapid Burster, and which was frequently observed to burst during targeted observations of the Rapid Burster (Fox et al. 2001).

2.2.4 Modelling the persistent emission

In order to measure a change in the level of persistent emission for every burst spectrum I need to construct a model for the persistent emission for each burst that can then be incorporated into the spectral fits. It is not possible to simply use the detected photon counts for each energy bin because these include instrumental background.

I found that the blackbody-plus-powerlaw model used for the example burst from 4U 1636–536 in §2.2 does not give adequate fits for every burst, so I considered a set of alternative spectral models. I consider any well-fitting phenomenological candidate of these models, as the main objective is only to get a function that matches the data, but

⁵see xte.mit.edu

the literature does provide some commonly-used functions. [White \(1986\)](#) points out the importance of scattering, mentioning that this can be modelled with a Compton spectrum or a blackbody with an added power law, and I include both in the selection of alternatives. The inclusion of a Gaussian at 6.4 keV was motivated by the detection of fluorescent Fe emission for some sources at this energy (G08). The combination of thermal bremsstrahlung with a Gaussian was found by experimenting in XSPEC with persistent spectra that could not be properly fit with any of the other five models. For models containing a Gaussian, the lower limit on the line width is set to 0.1 keV to avoid the Gaussian simply approximating a delta function that removes the error on a single energy bin. These are summarised in [Table 2.1](#), and I retained the fit that gave the best (i.e., lowest) χ^2_ν . I found that the six models listed in [Table 2.1](#) are sufficient to describe the persistent emission of all but one burst (see [Figure 2.6](#)). In models where the hydrogen column densities are held constant, these models are given in [Appendix A](#). The N_H values are mostly taken from the literature and references are given in [Appendix A](#). In persistent emission models where the hydrogen column densities are allowed to vary, the quantity $A(E)$ (see [equation 2.2.1](#)) is to be thought of as the product of the true interstellar absorption and a multiplicative factor that corrects the shape of the persistent emission model. As $A(E)$ is a multiplicative factor it is unaffected by changes in f_a and so this treatment does not introduce any systematic effects. One persistent emission spectrum, preceding a burst from Cyg X-2 observed at 14:29 March 27, 1996, (burst ID 2 in G08) could not be fit adequately by any of the six models, so this burst was excluded from the analysis.

The reduced χ^2 for the spectral fits to the remaining 332 persistent spectra had a mean of 1.03 and a standard deviation of 0.35, with an average of $\nu = 21$ degrees of freedom. The skewness was 1.00 and the kurtosis was 2.8. This distribution is therefore somewhat more skewed and significantly more peaked compared to expected theoretical values of $\sqrt{8/\nu} = 0.62$ and $12/\nu = 0.57$ for the skewness and kurtosis respectively. A Kolmogorov-Smirnov test gave a value of $D = 0.069$ and a 8.8% probability that the measured distribution is consistent with that expected assuming a good fit. I therefore consider these persistent emission models adequate for use in the subsequent generation of burst fits. The distribution of χ^2_ν for the fitted persistent emission spectra is shown in [Figure 2.6](#).

Table 2.1. XSPEC models for fitting the persistent emission

XSPEC model	Number of spectra	Notes
wabs*(bbodyrad+powerlaw)	127	nH fixed to values in Appendix A
wabs*(bbodyrad+powerlaw+gauss)	62	nH fixed, Gaussian energy set to 6.4keV
wabs*(compTT) ^a	37	nH fixed
wabs*(bbodyrad+powerlaw)	18	nH allowed to vary
wabs*(gauss+bremss)	31	all parameters variable ^b
wabs*(bbodyrad+diskbb) ^c	57	all parameters variable ^b
Total usable pre-burst spectra	332	
Rejected due to source confusion	36	Other active sources in field
No background data	3	
No good persistent model fit	1	Minimum $\chi^2_\nu > 5$
Total unusable bursts	40	

^aSee Titarchuk (1994a)

^bAll parameters are variable for the generation of persistent emission models. Their values are subsequently frozen for the burst spectral fits in §2.2.5 and Chapter 4

^cSee XSPEC manual (<http://heasarc.gsfc.nasa.gov/xanadu/xspec/manual/manual.html>) and references therein

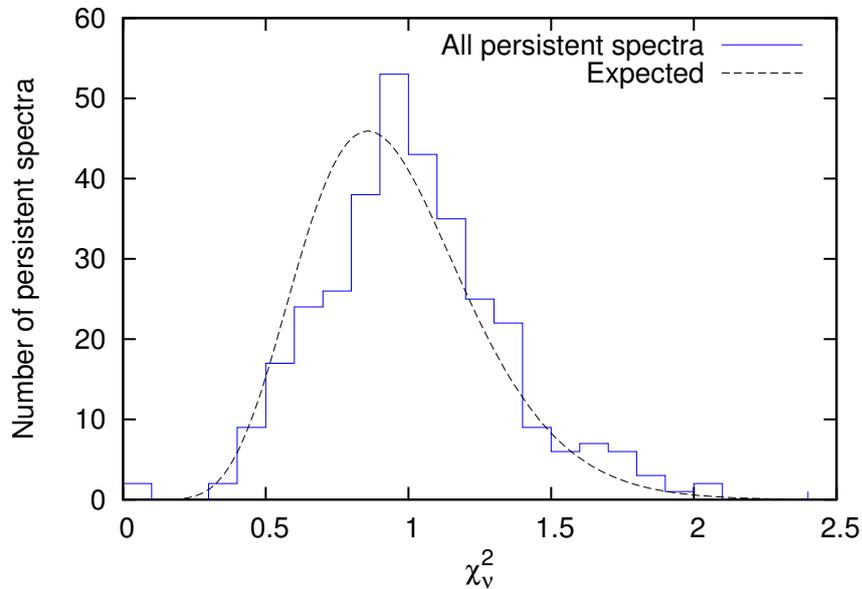


Figure 2.6 Distribution of χ^2_ν for the fits to the persistent emission spectra for the 332 PRE bursts (blue stepped line) and the theoretical distribution (dashed curve) for $\nu = 20$. A Kolmogorov-Smirnov test showed a 8.8% probability that the observed distribution is consistent with that expected assuming a good fit, indicating that the model fits of pre-burst spectra are adequate.

2.2.5 Fitting burst spectra

As in §2.2, initially I fit the burst spectra via the standard approach using an absorbed blackbody, by subtracting the measured pre-burst emission, i.e., `wabs*bbbodyrad` in XSPEC. This is the same implementation of the standard approach as the one outlined in G08.

I then replaced the detected pre-burst emission with just the instrumental background and fit the same spectra with a blackbody corrected by the adopted interstellar extinction value plus a multiple of the persistent emission model for that burst. I fix the hydrogen column density to the values given in Appendix A. The multiple of the persistent emission f_a is allowed to vary as a free parameter, along with the temperature and normalisation of the blackbody, as I did in §2.2.

The Levenberg-Marquardt algorithm used by XSPEC can become trapped in local χ^2 minima, returning obviously unphysical results such as extremely high temperatures or normalisations. This problem can be mitigated if sensible initial values for the parameters are given rather than the XSPEC defaults, which are not appropriate in every case. I use the blackbody temperature and normalisation of the blackbody obtained from the standard model, and $f_a = 1$, to seed these fits.

Generating and fitting to 1,000 simulated spectra, as was done in §2.2.2, for each of the tens of thousands of burst spectra would be computationally prohibitive. To overcome this problem, I excluded spectra for which the variable persistent flux fit was spurious, as follows. Only spectra with blackbody temperatures between 0.1 and 5.0 keV, normalisations of less than $10^6 \text{ km}^2/(10\text{km})^2$, and bolometric fluxes less than $10^{-6} \text{ erg cm}^{-2} \text{ s}^{-1}$ were retained. I also excluded f_a determinations for which XSPEC encountered fitting errors such as non-monotonicity or reaching the parameter limits, as described in the XSPEC manual⁶). This left 26,113 burst spectra out of the original 41,282. Almost all (>99%) of the discarded spectra were taken before the beginning of the burst or very late in the cooling tail, and the failure to obtain spectral fits can be attributed to very low photon counts for these spectra. I also reduced the number of simulated spectra for each measured spectrum from 1,000 to 320 to further reduce the computation time.

⁶see <http://heasarc.gsfc.nasa.gov/xanadu/xspec/manual/manual.htm>

Table 2.2. Stages of a PRE burst

Stage number	Stage name	Description
0	Pre-burst	Flux has not yet reached 1/4 of the maximum
1	Expansion	normalisation has not yet reached maximum
2	Contraction	Radius has reached maximum, kT has not yet reached maximum
3	Cooling tail	Flux still above 1/4 maximum
4	Post-burst	Flux has dropped below 1/4 maximum

As a further check that f_a is measuring a real spectral effect, I took the highest f_a spectrum for each burst and performed an f-test on it, comparing the χ^2 statistic from the standard approach and variable persistent flux fits. Because $f_a \times P(E)$ is an additive component (see equation 2.2.1), f-tests are a suitable test (e.g., [Orlandini et al. 2012](#)). I found that, of these, 165 had detections of greater than 3σ significance and 65 had detections of better than 4σ significance.

I define the *burst stage*: the *pre-burst stage* consists of all times before the beginning of the burst, the *expansion stage* occurs from the beginning of the burst to the time of maximum normalisation, the *contraction stage* is from the time of the maximum normalisation up to and including the touchdown time, and the *cooling tail stage* is all times after the touchdown time but before the bolometric flux drops back below one quarter of the maximum flux. I neglect spectra after this time. I refer to the expansion and contraction stages collectively as the *Eddington-limited stage*. These stages are summarised in Table 2.2.

2.3 Eddington Fluxes

I also give Eddington fluxes, calculated from parameters measured with the standard approach, for all sources that have PRE bursts observed by either *RXTE* (i.e. in the extended G08 catalogue) or the Wide Field Camera on *BeppoSAX*⁷. The Eddington fluxes are the means of the peak fluxes for every PRE burst detected for that source, weighted by the inverse square of the error of the measurement. These are listed in Table A.1, in Appendix A.

⁷see burst.sci.monash.edu/minbar

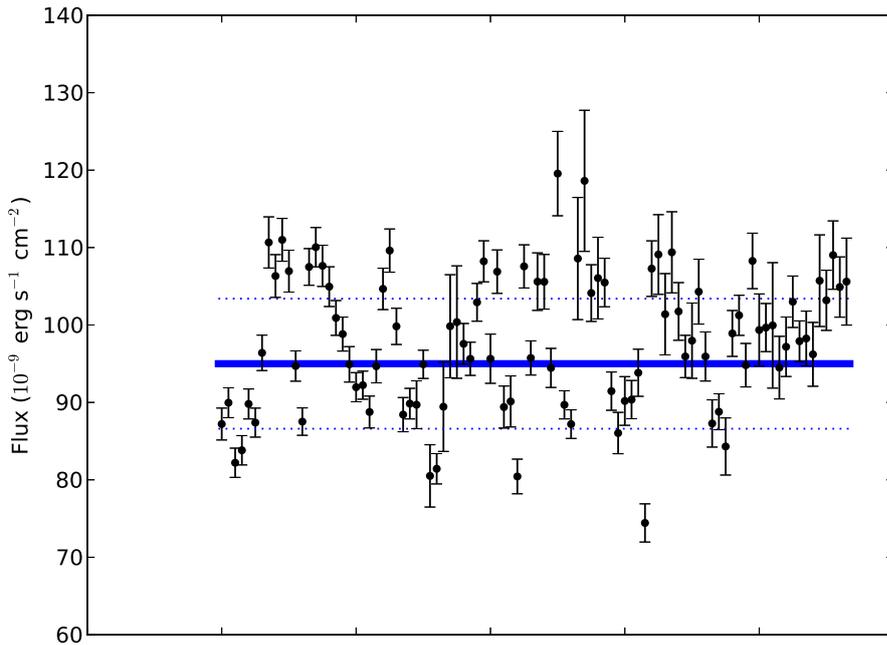


Figure 2.7 Measurement of the Eddington flux of 4U 1728–34 from the peak fluxes of its individual radius expansion bursts. The thick blue line is the calculated Eddington flux and the dotted blue lines are the $\pm 1\sigma$ uncertainties on that measurement. All radius expansion bursts for this source were detected by *RXTE*.

I calculated the Eddington flux for each source as follows: if only one PRE burst has been recorded for any source, I simply reported the peak flux for that burst and the error on that measurement. For sources with two or more PRE bursts I calculated a reduced χ_{meas}^2 assuming the source has constant Eddington flux. If χ_{meas}^2 is consistent with a constant flux, I allow the flux to vary up and down such that it now has $\chi_{\text{new}}^2 = 1 + \chi_{\text{meas}}^2$; the amount by which I varied it is the error on the original measurement, where the subscripts “meas” and “new” indicate the values initially measured and the newly adjusted values respectively. If χ_{meas}^2 is *not* consistent with constant Eddington flux I artificially scaled the errors on the individual measurements until $\chi_{\text{meas}}^2 = 1$. Then I allow the flux to vary such that $\chi_{\text{new}}^2 = 2$. There are more sophisticated ways of combining seemingly inconsistent measurements of a quantity (e.g., [Press 1996](#)), but this method is sufficient for the purpose. Figure 2.7 illustrates the method for the burster 4U 1728–34.

The prolific burster 4U 1636–536 has PRE bursts that are known to be bimodal in flux ([Ebisuzaki & Nakamura 1988](#), [Galloway et al. 2006](#)). Most of its PRE bursts are thought to take place in a pure helium atmosphere and have peak fluxes around $(68.6 \pm 5.5) \times 10^{-9} \text{ erg cm}^{-2} \text{ s}^{-1}$. A few bursts have peak fluxes lower by a factor of ≈ 1.7 , and these

are believed to take place in a hydrogen-rich atmosphere (G08). I consider both regimes separately. The lower flux bursts include a tentatively identified (i.e., marginal) PRE event observed by *RXTE* on 2000 Jan 22, 04:43:48 UT. I include this burst to calculate the mean Eddington flux for these low flux events, but exclude it for the rest of the analyses in this thesis.

Chapter 3

SPH Modelling

As soon as we started programming, we found to our surprise that it wasn't as easy to get programs right as we had thought. Debugging had to be discovered. I can remember the exact instant when I realized that a large part of my life from then on was going to be spent in finding mistakes in my own programs.

Maurice Wilkes

I aim to use three-dimensional computer simulations to test the hypothesis that radiation drag during a type I X-ray burst can increase the accretion rate. In this chapter I introduce the equations of fluid dynamics which govern the behaviour of the accretion disc, including a velocity-dependent term representing radiation drag. I review the Smoothed Particle Hydrodynamics (SPH) technique I will use.

3.1 Fluid dynamics

3.1.1 Conservation of mass

Suppose I have a quantity of gas characterised by its density ρ , temperature T , and velocity \mathbf{v} , these three quantities being functions of position \mathbf{r} and time t . The behaviour of this gas is described by three conservation equations— those of mass, momentum, and energy.

The *continuity equation* states that the rate of change of density inside a volume element is the rate at which material flows into or out of it:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{v}) = S, \quad (3.1.1)$$

where S allows for the possibility of *sources* or *sinks* which, respectively, create or destroy gas. Usually S will be zero everywhere, but if accreting objects are present, they can be described as locations where S is negative. Equation 3.1.1 describes the rate of change of density with respect to a fixed coordinate system; the same can be expressed for a coordinate system moving with the gas (i.e., the Lagrangian) as

$$\frac{D\rho}{Dt} = -\rho \nabla \cdot \mathbf{v}, \quad (3.1.2)$$

which follows directly from the definition of the moving derivative of a general quantity ψ

$$\frac{D\psi}{Dt} = \frac{\partial \psi}{\partial t} + (\nabla \psi) \cdot \mathbf{v}. \quad (3.1.3)$$

As I will explain later, moving derivatives are natural for SPH because in quantities and derivatives are calculated on discrete moving particles.

3.1.2 Conservation of momentum

Pressure gradients and external forces drive motions of material within the gas, and viscous effects partially oppose them. These processes are described by the *Navier-Stokes equation*, which I give here in co-moving form:

$$\frac{Dv^i}{Dt} = \frac{1}{\rho} \frac{\partial T^{ij}}{\partial x^j} + f_{\text{ext}}^i, \quad (3.1.4)$$

where the force density \mathbf{f} quantifies other forces, such as gravity, magnetic fields, and radiation force acting upon the gas. The stress tensor T contains terms corresponding to pressure and viscosity within the gas.

The stress tensor is given by

$$T^{ij} = \left[-P + \left(\zeta - \frac{2}{3}\eta \right) \frac{\partial v^k}{\partial x^k} \right] \delta^{ij} + \eta \left(\frac{\partial v^i}{\partial x^j} + \frac{\partial v^j}{\partial x^i} \right). \quad (3.1.5)$$

In this expression P is the pressure, and $\zeta = \nu\rho$ and η are the shear and bulk components of the viscosity respectively. As mentioned in Chapter 1, the bulk viscosity is (or should be) negligible because it is dynamically important only for episodes of very rapid expansion or compression.

The gas pressure P is a consequence of gas particles colliding through thermal motions. It is directly proportional to density and temperature:

$$P = \frac{\rho k T}{\mu m_H}, \quad (3.1.6)$$

where k is the Boltzmann constant and m_H is the mass of hydrogen. The mean molecular weight μ is the mean weight of the gas particles in atomic mass units. I assume the gas is composed purely of hydrogen and helium, with other species only present in trace amounts small enough to be ignored. For atomic hydrogen $\mu = 1$ whereas ionised hydrogen has $\mu = \frac{1}{2}$. Ionised material of solar composition (75% hydrogen, 25% helium) has $\mu = 0.59$. In general, if the hydrogen mass fraction is X , then

$$\mu = \frac{4}{3 + 5X}. \quad (3.1.7)$$

3.1.3 Conservation of Energy

The final equation is conservation of energy. It is similar in form to Equation 3.1.1, except that the conserved quantity is energy (both kinetic and internal), and there are several source and sink terms. In a co-moving coordinate system it is:

$$\frac{Du}{Dt} = -\frac{P}{\rho} (\nabla \cdot \mathbf{v}) - \nabla \cdot \mathbf{q} - \nabla \cdot \mathbf{F}_{\text{Rad}} \quad (3.1.8)$$

The $\nabla \cdot \mathbf{F}_{\text{Rad}}$ term of 3.1.8 describes the gain or loss of energy through the absorption and emission of radiation, and \mathbf{q} is the conductive heat flux. Temperature gradients are usually small enough that \mathbf{q} can be neglected.

Given descriptions of \mathbf{f} and \mathbf{F}_{Rad} , equations 3.1.2, 3.1.4, and 3.1.8 provide a means of modelling the behaviour of the gas. In practice, however, the radiation and term complicates the system considerably, because to treat it properly entails calculating heating from viscosity, from the star, and from the disc itself. Doing this properly involves computationally expensive ray-tracing. Cooling of the disc also needs to be taken into account. It is usual to simply matters by neglecting \mathbf{F}_{Rad} as well. It then follows that

$$\frac{Du}{Dt} = \frac{P}{\rho^2} \frac{D\rho}{Dt}, \quad (3.1.9)$$

from which

$$\begin{aligned} u &= K\rho^{\gamma-1} \\ P &= K\rho^\gamma \end{aligned} \quad (3.1.10)$$

for some constants γ, K .

Another important expression for the pressure, valid when temperature is locally constant (i.e., $\gamma = 1$), is

$$P = c_s^2(R)\rho, \quad (3.1.11)$$

where c_s is the local speed of sound. That is,

$$c_s^2(R) = \frac{kT(R)}{\mu m_h}. \quad (3.1.12)$$

The sound speed will be used in descriptions of the accretion disc structure (see §1.6). Note that the temperature can, in general, vary as a function of radial distance from the star. However, it is assumed not to vary in the vertical or azimuthal directions, and its radial gradient is assumed to be small. These assumptions are referred to as *locally isothermal*.

3.2 Formulation in Smoothed Particle Hydrodynamics

3.2.1 Overview of SPH

Smoothed particle hydrodynamics (SPH) is a Lagrangian method for solving the equations of hydrodynamics in which the fluid is discretised into N particles, each carrying a fixed mass, that follow the fluid motion and on which the fluid quantities and their derivatives are computed (Lucy 1977, Gingold & Monaghan 1977; see Monaghan 2005, Price 2012 for reviews). It has found widespread use in a variety of scientific disciplines, including astrophysics, geophysics and engineering, and has also been applied to computer graphics.

An advantage over grid-based methods is that SPH adapts naturally to changes in the shape of the distribution and flow of mass through the motion of the SPH particles; simply put, the resolution is finest wherever all the mass is. SPH does not have external boundaries, which means that no mass can escape the simulation and simulations are not subject to artificial restrictions on the geometry of the systems being studied. Derivatives are calculated by interpolating over nearby particles, adding considerable computational cost because the neighbours of the particles need to be identified frequently.

I describe the basis of SPH here, following Monaghan (1992). Consider some arbitrary variable $A(\mathbf{r})$ that could be associated with a particle. It is trivially true that

$$A(\mathbf{r}) = \int A(\mathbf{r}')\delta(|\mathbf{r} - \mathbf{r}'|) d\mathbf{r}', \quad (3.2.1)$$

where δ is Dirac's delta function and the integral is performed over all of space. The idea behind SPH is to replace the δ function with a *kernel*, $W(|\mathbf{r} - \mathbf{r}'|, h)$, with a characteristic width h and normalised so that

$$\int W(\mathbf{r} - \mathbf{r}', h) d\mathbf{r}' = 1. \quad (3.2.2)$$

It is also necessary that the kernel approaches the δ function as h approaches zero, so that nearby particles affect the evaluations more than distant ones. I still have

$$A(\mathbf{r}) = \int \frac{A(\mathbf{r}')}{\rho(\mathbf{r}')} W(|\mathbf{r} - \mathbf{r}'|, h) \rho(\mathbf{r}') d\mathbf{r}' + O(h^2), \quad (3.2.3)$$

where the error term is quadratic if the kernel is chosen to be an even function of \mathbf{r} , which I will henceforth assume.

Now I may discretise this integral onto the particles by replacing it with a summation over all the particles, and recognising that the mass element ρdV is just the particle mass m . Now,

$$A(\mathbf{r}) \approx \sum_{b=1}^N m_b \frac{A_b}{\rho_b} W(|\mathbf{r} - \mathbf{r}'|, h), \quad (3.2.4)$$

where quantities subscripted b are evaluated at the location of particle b . Any quantity of the fluid under consideration (density, pressure, temperature, etc.) can take the place of A in this above expression. The density itself is, by trivial cancellation of ρ :

$$\rho(\mathbf{r}) \approx \sum_{b=1}^N m_b W(|\mathbf{r} - \mathbf{r}'|, h). \quad (3.2.5)$$

Furthermore, derivatives of A can be expressed as summations over the derivatives of the kernel:

$$\frac{\partial A}{\partial x} = \sum_{b=1}^N m_b \frac{A_b}{\rho_b} \frac{\partial W}{\partial x}. \quad (3.2.6)$$

In practice, however, better accuracy is achieved (Monaghan 1992) if we use the representation

$$\rho \nabla A = \nabla(\rho A) - A \nabla \rho, \quad (3.2.7)$$

so that ∇A becomes

$$\nabla A(\mathbf{r}) = \frac{1}{\rho} \sum_{b=1}^N m_b (A(\mathbf{r}) - A(\mathbf{r}_b)) \nabla W(|\mathbf{r} - \mathbf{r}_b|, h). \quad (3.2.8)$$

Similarly, $\nabla \cdot \mathbf{v}$ is better written as

$$\nabla \cdot \mathbf{v} = \frac{1}{\rho} \sum_{b=1}^N m_b (\mathbf{v} - \mathbf{v}_b) \cdot \nabla W(|\mathbf{r} - \mathbf{r}_b|, h). \quad (3.2.9)$$

3.2.2 Kernel choice

Selecting an appropriate functional form for the kernel is of fundamental importance. An obvious choice of kernel is a Gaussian normalised to unity, but this has the disadvantage

that it is computationally expensive. Furthermore, it is nonzero everywhere, which means either summing over every particle or truncating the Gaussian at some radius, leading to discontinuities. It is better to choose a simple polynomial, or splines composed of polynomial segments. One possibility is

$$M_2(x; h) = \begin{cases} 1 - q, & q < 1 \\ 0, & q \geq 1, \end{cases} \quad (3.2.10)$$

where $q = |x|/h$. This is very fast computationally, but has the drawback that its derivative is discontinuous. Furthermore, it is better for the density estimate to be insensitive to small changes in the position of a near neighbour (Price 2012). This criterion is only achievable if the kernel is approximately flat around the origin, but it is not true for the linear kernel. The favoured kernel is a cubic spline, which is Gaussian-like in shape but has finite extent (i.e., it has *compact support*):

$$M_4(x; h) = \begin{cases} k [(2 - q)^3 - 4(1 - q)^3], & 0 \leq q \leq 1 \\ k(2 - q)^3, & 1 \leq q \leq 2 \\ 0, & q \geq 2, \end{cases} \quad (3.2.11)$$

where $k = [\frac{1}{6h}, \frac{5}{14\pi h^2}, \frac{1}{4\pi h^3}]$ in $[1, 2, 3]$ dimensions, which has continuous first and second derivatives everywhere, and is zero outside $q \geq 2$. This means that the algorithm need not sum over every particle, but only nearby ones. Since kernels are computed frequently, it is customary to precompute tables of values for $W(q)$ and $\partial W/\partial q$, and find intermediate values by linear interpolation. This leads to the same computational time regardless of the functional form of the kernel. The Gaussian, M_2 , and M_4 kernels and their first derivatives are shown in Figure 3.1. There are a large number of other possible kernels, but there is no consensus on a better default than the m_4 kernel for similar computational cost (e.g. Price 2012).

3.2.3 Smoothing length

The characteristic width, or *smoothing length*, h specifies the rate at which the kernel falls off with distance and is, in general, a function of location and particle spacing. It is

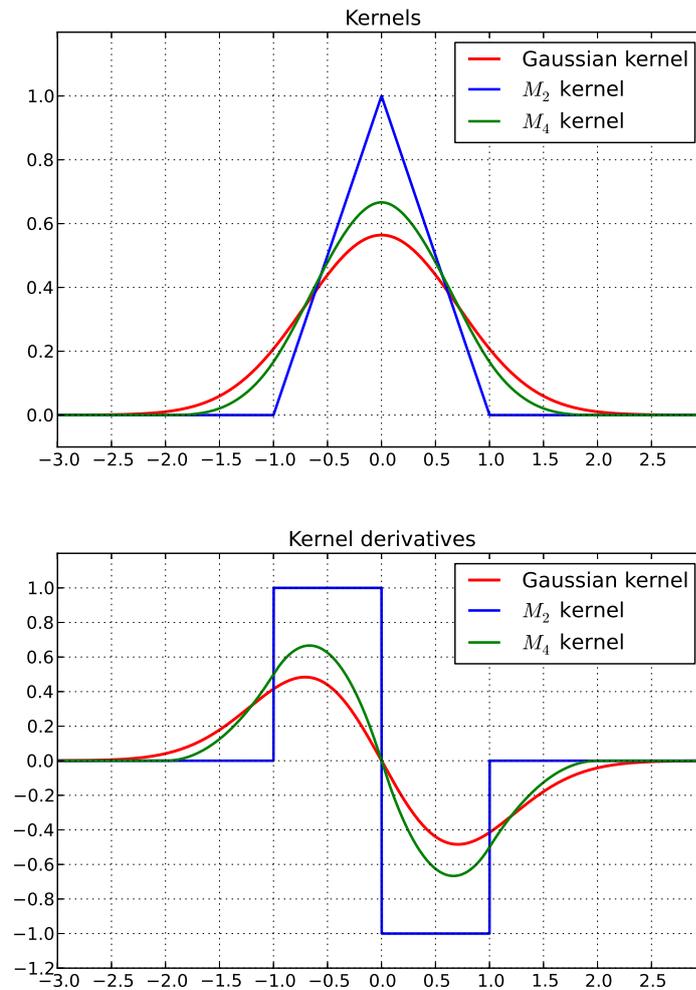


Figure 3.1 Three choices of kernels, with $h = 1$, and their derivatives for one dimension and $h = 1$. The Gaussian kernel is infinitely differentiable, but nonzero everywhere. The linear spline M_2 is fast to compute and is zero for $q \geq 1$ but it does not have continuous derivatives. M_4 has compact support and continuous first and second derivatives and, being a cubic, is not prohibitively expensive. These properties make it a popular choice for SPH kernels.

desirable to have h small where particles are close together and large where the density is low, so that the smoothing length is proportional to the average particle spacing. For particles of equal mass, this is the same as making $h(\mathbf{r})$ proportional to $\rho(\mathbf{r})^{-1/d}$, where d is the number of dimensions. But since the density is a function of smoothing length (see equation 3.2.5) we end up with a set of simultaneous equations for h_a and ρ_a at every particle a :

$$\begin{aligned} \rho(\mathbf{r}_a) &= \sum_{b=1}^N m_b W_{ab}, \text{ and} \\ h(\mathbf{r}_a) &= \eta \left(\frac{m_a}{\rho_a} \right)^{1/d}, \end{aligned} \tag{3.2.12}$$

where η is a parameter specifying the smoothing length in units of the mean particle spacing $(m/\rho)^{1/d}$. I have also introduced the shorthand W_{ab} to stand for $W(|\mathbf{r}_a - \mathbf{r}_b|, h_a)$. These simultaneous equations are solved for every particle with standard root-finding methods, such as the Newton-Raphson method. Another method is to set h so that all particles have the same number of neighbours (Hernquist & Katz 1989).

3.2.4 Error estimates

Errors in the calculated quantities and their derivatives are not easy to estimate for SPH. They should decrease with increasing N , but they do not decrease as fast as N^{-1} due to the tendency of SPH particles to move themselves into partially ordered arrangements (Monaghan 2005, Price 2012). The reason for this is intuitively obvious. Consider a region of uniform density. If the particles were randomly distributed throughout this region, there would be some areas of localised over- or under-density because there will exist particles whose nearest neighbour is much closer or much further than the average particle spacing. When the gas evolves such density anomalies will be smoothed out, leading to a much more uniform particle spacing in the form of a partially ordered arrangement of particles (see Figure 3.2). The true decrease in errors with particle number in practice seems to be between N^{-1} and $N^{-1/2}$ (e.g., Thompson 2006) and it can be shown that the best achievable is $N^{-2/3}$ (Price 2012).

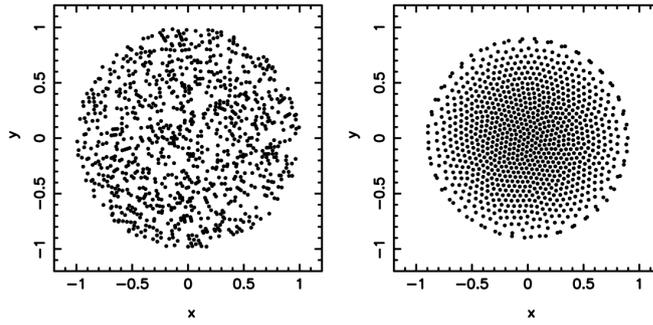


Figure 3.2 Illustration of particle self-ordering from [Monaghan \(2005\)](#). The SPH particles were initially placed randomly within a unit sphere (left panel), then allowed to evolve under a gravitational potential (right panel). The ordered structure that the particles reach is evident in the right panel.

3.3 Fluid equations in SPH

In this work I use the PHANTOM code developed by [Price \(2014\)](#). This is a highly efficient SPH code optimised for studying non-self-gravitating problems. As shown in previous work (e.g., [Lodato & Price 2010](#)) PHANTOM is suitable for studying viscous discs evolving under a gravitational potential and velocity dependent torques. Since Poynting-Robertson drag is also a velocity dependent torque, I expect that PHANTOM will be able to deal with the irradiated neutron star accretion discs under discussion.

3.3.1 Conservation of mass in SPH

The SPH method has long been known to be suitable for the study of viscous accretion discs ([Artymowicz & Lubow 1994](#), [Murray 1996](#)). Its formulation ensures that the basic equations of hydrodynamics are elegant and simple in SPH. For example, consider the SPH expression for density:

$$\rho_a = \sum_{b=1}^N m_b W_{ab}. \quad (3.3.1)$$

If I now take the Lagrangian (i.e. co-moving) derivative I obtain

$$\frac{D\rho_a}{Dt} = \sum_{b=1}^N m_b (\mathbf{v}_a - \mathbf{v}_b) \cdot \nabla_a W_{ab}. \quad (3.3.2)$$

But this is just the discretised version of the continuity equation

$$\frac{D\rho}{Dt} = -\rho(\nabla \cdot \mathbf{v}). \quad (3.3.3)$$

We therefore see that the continuity equation has a very simple and useful expression in SPH as summations over particles.

3.3.2 Conservation of momentum in SPH

The Navier-Stokes equation in SPH, for particle a , is

$$\frac{Dv_a^i}{Dt} = \frac{1}{\rho_a} \frac{\partial T_a^{ij}}{\partial x^j} + f_{a,\text{ext}}^i. \quad (3.3.4)$$

In SPH it is customary to add an extra term, the *artificial viscosity* (q^{AV}), to the stress tensor so that we have

$$T_a^{ij} = \left[-(P + q_a^{AV}) + \left(\zeta_a - \frac{2}{3}\eta_a \right) \frac{\partial v_a^k}{\partial x_a^k} \right] \delta^{ij} + \eta_a \left(\frac{\partial v_a^i}{\partial x_a^j} + \frac{\partial v_a^j}{\partial x_a^i} \right). \quad (3.3.5)$$

Here and henceforth the superscript AV stands for artificial viscosity. Note that the bulk and shear viscosities can vary between particles, as emphasised by the subscript a . The purpose of artificial viscosity is to handle shocks, and also to prevent SPH particles piling up on top of one another. It is possible to interpret q^{AV} as a physical viscosity, by setting η and ζ to zero, but the drawback is that one can no longer adjust the shear and bulk components individually. In systems like accretion discs, which should not have any bulk viscosity, this is a disadvantage. On the other hand, it means that viscosity can be parameterised by a single viscosity coefficient. Furthermore, having a large unphysical bulk viscosity turns out not to be such a problem because $\nabla \cdot \mathbf{v}$ is generally small.

The artificial viscosity in PHANTOM is given by

$$q_a^{AV} = \frac{1}{2} \alpha_a^{AV} \rho_a v_{\text{sig},a} |\mathbf{v}_{ab} \cdot \hat{\mathbf{r}}_{ab}| \quad (3.3.6)$$

if $\mathbf{v}_{ab} \cdot \hat{\mathbf{r}}_{ab}$ is negative, and zero otherwise. That is, it applies only to converging flows. The *signal velocity* is

$$v_{\text{sig},a} = c_{s,a} + 2 |\mathbf{v}_{ab} \cdot \hat{\mathbf{r}}_{ab}|, \quad (3.3.7)$$

with the final term providing a strong resistance to motion in fast shocks to prevent interpenetration of particles. In Equation 3.3.7 α^{AV} is the artificial viscosity. If, as discussed above, we let ζ and η be zero, all of the viscous effects are handled with α^{AV} .

The artificial viscosity can be related to the Shakura-Sunyaev viscosity α by making three minor adjustments to Equation 3.3.6, as described by Lodato & Price (2010):

- $v_{\text{sig},a} = c_{s,a}$, i.e., the shock resistant term is removed.
- q_a^{AV} is multiplied by $h/|\mathbf{r}_{ab}|$.
- The restriction to convergent flows is removed.

It can then be shown that

$$\alpha \approx \frac{1}{10} \alpha^{AV} \frac{\langle h \rangle}{H}, \quad (3.3.8)$$

where $\langle h \rangle$ is the average SPH particle smoothing length for a given radius in the disc; see (Lodato & Price 2010) for the derivation. To ensure a constant α throughout a disc it is necessary to set the disc up in such a way that $\langle h \rangle / H$ is constant.

3.3.3 Timestepping in PHANTOM

PHANTOM uses the Velocity Verlet approach. A particle at position \mathbf{r}^n time step n is moved to its new position \mathbf{r}^{n+1} at the next timestep according to

$$\mathbf{r}^{n+1} = \mathbf{r}^n + \Delta t \mathbf{v}^n + \frac{1}{2} \Delta t^2 \mathbf{a}^n, \quad (3.3.9)$$

where $\mathbf{a} = \mathbf{f}/m$ is the acceleration of the particle. Observe that this can be written as

$$\begin{aligned} \mathbf{r}^{n+1} &= \mathbf{r}^n + \Delta t \left[\mathbf{v}^n + \frac{1}{2} \Delta t \mathbf{a}^n \right] \\ &= \mathbf{r}^n + \Delta t \mathbf{v}^{n+1/2}, \end{aligned} \quad (3.3.10)$$

where $\mathbf{v}^{n+1/2}$ is to be thought of as the velocity at a half-timestep. After this step, \mathbf{a} is re-evaluated. Thus, the velocity is evaluated from position and acceleration at full timesteps, and position is evaluated from velocities at half timesteps, i.e.,

$$\mathbf{v}^{n+3/2} = \mathbf{v}^{n+1/2} + \frac{1}{2}\Delta t [\mathbf{a}^n + \mathbf{a}^{n+1}]. \quad (3.3.11)$$

The method works provided acceleration depends only upon \mathbf{r} and not upon \mathbf{v} . In this work acceleration *is* a function of velocity because of radiation drag. In §3.4 I explain how this issue is dealt with.

3.4 The extern_prdrag module

The `extern_prdrag` module specifies the external force due to a central point-like, gravitating, emitting, and non-rotating object. It is intended to be as general as possible. The equation of motion of a particle subject to these forces is

$$\mathbf{a}_{\text{ext}} = -\frac{(1 - k_0\beta)GM_*}{r^3}\mathbf{r} - \beta \left(k_1 \frac{GM_*}{r^3} \frac{v_r}{c} \mathbf{r} - k_2 \frac{GM_*}{r^2} \frac{\mathbf{v}}{c} \right), \quad (3.4.1)$$

where \mathbf{r} and \mathbf{v} are position and velocity, v and r are their magnitudes, and v_r is the component of the velocity in the radial direction. M_* is the mass of the central object. The acceleration \mathbf{a} is, of course, the force on an SPH particle divided by the mass of the particle. The parameter β is the ratio of radiation to gravitational forces, and the parameters k_0, k_1, k_2 are input file options, which I will describe later. All other symbols have their customary meanings. Since the emitting object is assumed to be a point, small corrections associated with an extended emitter (Guess 1962) are neglected. Relativistic effects are likewise neglected, as these are thought to be less important than radiation forces even for very low ($\beta < 0.01$) stellar luminosities (Miller & Lamb 1993).

The first terms on the right side of Equation 3.4.1 corresponds to gravity, reduced by outward radiation pressure. The last term represents drag against the direction of motion of the particle, and the middle term is a redshift term. The `extern_prdrag` module allows these three terms to be scaled independently through the three input file parameters k_0, k_1, k_2 , whose defaults are unity– they are called `RadiationPressure`, `Redshift`, and

`TransverseDrag` in the input file. Rotation of the neutron star, which is not included in this work, can be crudely emulated by changing k_2 .

A difficulty with the Velocity Verlet scheme for updating the velocity is that the updated velocity \mathbf{v}^1 depends upon the updated acceleration \mathbf{a}^1 :

$$\mathbf{v}^1 = \mathbf{v}^0 + \frac{1}{2}\Delta t (\mathbf{a}^0 + \mathbf{a}^1), \quad (3.4.2)$$

where the superscripts 1,0 indicate the old and updated values of \mathbf{v} and \mathbf{a} . However, \mathbf{a}^1 depends upon \mathbf{v}^1 and so I must solve implicitly. If I expand the \mathbf{a}^1 term of the leapfrog prescription into velocity-dependent and non velocity-dependent components I have

$$\begin{aligned} \mathbf{v}^1 &= \mathbf{v}^0 + \frac{1}{2}\Delta t \mathbf{f}^0 + \frac{1}{2}\Delta t \mathbf{a}_{nvd}^1 + \frac{1}{2}\Delta t \mathbf{a}_{vd}^1 \\ &= \mathbf{T} + \frac{1}{2}\Delta t \mathbf{a}_{vd}^1, \end{aligned} \quad (3.4.3)$$

where

$$\begin{aligned} \mathbf{T} &= \mathbf{v}^0 + \frac{1}{2}\Delta t \mathbf{a}^0 + \frac{1}{2}\Delta t \mathbf{a}_{nvd}^1 \\ &= \mathbf{v}^0 + \frac{1}{2}\Delta t \mathbf{a}^0 - \frac{(1 - k_0\beta)GM_*\Delta t}{2r^3} \mathbf{r} \end{aligned} \quad (3.4.4)$$

and consists only of the velocity and force from the previous step (therefore already calculated) and radial force (which does not involve \mathbf{v}^1). Now I consider the velocity dependent part:

$$\begin{aligned} \frac{1}{2}\Delta t \mathbf{f}_{vd}^1 &= -\frac{1}{2}\Delta t \beta \left(k_2 \frac{GM_*}{r^3} \frac{v_r^1}{c} \mathbf{r} - k_1 \frac{GM_*}{r^2} \frac{\mathbf{v}^1}{c} \right) \\ &= -\frac{GM_*\beta\Delta t}{2cr^2} (k_2 \mathbf{v}^1 + k_1 v_r^1 \hat{\mathbf{r}}) \\ &= -Q (k_2 \mathbf{v}^1 + k_1 v_r^1 \hat{\mathbf{r}}), \end{aligned} \quad (3.4.5)$$

where

$$Q = \frac{GM_*\beta\Delta t}{2cr^2}. \quad (3.4.6)$$

Substituting into Equation 3.4.3, the leapfrog scheme becomes

$$\mathbf{v}^1 = -\frac{\mathbf{T} - Qk_1 v_r^1 \hat{\mathbf{r}}}{1 + Qk_2}. \quad (3.4.7)$$

If it is assumed that $|v_r| \ll |v|$ then this simplifies to

$$\mathbf{v}^1 \approx -\frac{\mathbf{T}}{1 + Qk_2}. \quad (3.4.8)$$

If not, then I use $v_r = \mathbf{v} \cdot \hat{\mathbf{r}}$ to obtain:

$$\mathbf{v}^1 = -\frac{\mathbf{T} - Qk_1(\mathbf{v}^1 \cdot \hat{\mathbf{r}})\hat{\mathbf{r}}}{1 + Qk_2}, \quad (3.4.9)$$

which, after some rearranging, gives a set of simultaneous equations for the components of \mathbf{v} :

$$\begin{aligned} v_x [Qk_1x^2 - (1 + Qk_2)r^2] &= T_x r^2 - Qk_1x(yv_y + zv_z) \\ v_y [Qk_1y^2 - (1 + Qk_2)r^2] &= T_y r^2 - Qk_1y(xv_x + zv_z) \\ v_z [Qk_1z^2 - (1 + Qk_2)r^2] &= T_z r^2 - Qk_1z(xv_x + yv_y). \end{aligned} \quad (3.4.10)$$

This system can be solved analytically:

$$\begin{aligned} -Sv_x &= (-T_y y - T_z z)Qk_1x + Q(k_1 - k_2)T_x r^2 - T_x (r^2 + Qk_1x^2), \\ -Sv_y &= (-T_x x - T_z z)Qk_1y + Q(k_1 - k_2)T_y r^2 - T_y (r^2 + Qk_1y^2), \\ -Sv_z &= (-T_x x - T_y y)Qk_1z + Q(k_1 - k_2)T_z r^2 - T_z (r^2 + Qk_1z^2), \end{aligned} \quad (3.4.11)$$

where

$$S = [(k_1k_2 - k_2^2)Q^2 + (k_1 - 2k_2)Q - 1]r^2. \quad (3.4.12)$$

If $k_1 = k_2 = 1$, under default conditions, there is convenient cancellation and the system reduces to

$$\begin{aligned} Sv_x &= (-T_y y - T_z z)Qx - T_x (r^2 + Qx^2), \\ Sv_y &= (-T_x x - T_z z)Qy - T_y (r^2 + Qy^2), \\ Sv_z &= (-T_x x - T_y y)Qz - T_z (r^2 + Qz^2), \end{aligned} \quad (3.4.13)$$

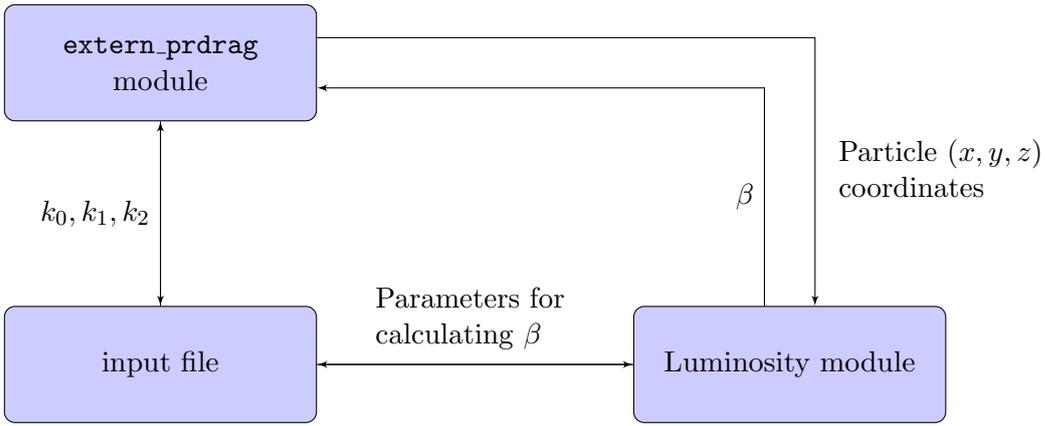


Figure 3.3 Relationships between the `extern_prdrag` module, the PHANTOM input file, and the user-specified luminosity module. The `extern_prdrag` module supplies only the spatial coordinates of the particles to the luminosity module, which returns only β . None of the details of calculating β are known to `extern_prdrag`. Both modules supply and read their own input parameters to the input file but, again, these are not communicated between the two modules.

with

$$S = (Q + 1)r^2. \quad (3.4.14)$$

We can substitute these back into equation 3.4.3 to get the velocity dependent part of the force.

The `extern_prdrag` module is designed to be as general as possible, with users free to prescribe β as they desire. This can be done by writing a new module that returns β to `extern_prdrag`. The new module should be able to write any changeable physical parameters to the input file, but `extern_prdrag` will not refer to them in any way (see Figure 3.3).

3.4.1 Description of Poynting-Robertson drag parameter β

Poynting-Robertson drag is usually formulated in the context of discrete particles with surface areas and masses (e.g., dust grains or rocks). Here I provide a description for a continuous absorbing medium, relevant for accumulations of gas such as accretion disks. Recall from §3.4 that the equation of motion of a particle under the influence of gravity and radiation drag is

$$\frac{d^2\mathbf{r}}{dt^2} = -\frac{(1-\beta)GM_*}{r^3}\mathbf{r} - \beta\left(\frac{GM_*}{r^3}\frac{v_r}{c}\mathbf{r} + \frac{GM_*}{r^2}\frac{\mathbf{v}}{c}\right), \quad (3.4.1)$$

where β is the ratio of the radiation force to the gravitational force. This description can be used for SPH particles representing a portion of disc gas as well as physical particles.

We begin with the definition of β :

$$\beta = \frac{F_{\text{Rad}}}{F_{\text{Grav}}}. \quad (3.4.15)$$

Now consider an element of the accretion disk with density ρ , surface area dA pointed toward the star, and radial extent dr . Suppose also that the disk material has constant, frequency independent, opacity κ , and let the flux at the location of the disk element be \mathcal{F} .

Obviously

$$F_{\text{Grav}} = \frac{GM_*\rho dr dA}{r^2}, \quad (3.4.16)$$

where M_* is the mass of the neutron star and r is the distance from the disk element to the center of the neutron star.

Now for the derivation of the radiation force. I have a flux of \mathcal{F} arriving at the disk element and, by the Beer-Lambert law, $\mathcal{F} \exp[-\rho\kappa dr]$ transmitted through it. Thus, an amount $\mathcal{F}(1 - \exp[-\rho\kappa dr])$ has been absorbed by the disk element. Let us take the Maclaurin series of the exponential function and neglect higher powers of dr . It follows that the flux absorbed by the disk element is $\mathcal{F}\rho\kappa dr$.

But $\mathcal{F} = L_{\text{local}}dA/c$, where L_{local} is the luminosity arriving at the element, so I get the radiation force on the disk element:

$$F_{\text{Rad}} = \frac{\kappa L_{\text{local}}\rho dr dA}{4\pi cr^2}. \quad (3.4.17)$$

Substituting into equation 3.4.15,

$$\beta = \frac{\kappa L_{\text{local}}}{4\pi GcM_*} = \frac{L_{\text{local}}}{L_{\text{Edd}}}. \quad (3.4.18)$$

Note that general relativistic effects do not apply to β because both the local and Eddington luminosity vary by the same factor. Thus the β parameter is simply the ratio of the local to the Eddington luminosity. It is identical to the quantity Γ considered by, e.g., [Abramowicz et al. \(1990\)](#). The study of [W92](#) uses luminosity rather than β , but β can be easily obtained by dividing by the known Eddington luminosity of those simulations.

Chapter 4

Observational Results: Radius Expansion Bursts

I didn't think it was physically possible, but this both sucks and blows.

Bart Simpson

In this Chapter, I have applied the method developed in Chapter 2 to the entire sample of PRE bursts detected by *RXTE*.

4.1 Distribution of f_a

I compared the distributions of χ^2_ν obtained with the variable flux approach to the expected distribution assuming a correct model using the Kolmogorov-Smirnov test. For the cooling tail I obtained $D = 0.304$ for the standard approach fit and $D = 0.193$ for the variable fit. For the Eddington-limited spectra I obtained $D = 0.331$ and $D = 0.171$ respectively. These results indicate that allowing f_a to vary consistently improves the goodness-of-fit, and the improvement is more pronounced for the Eddington-limited spectra than for the cooling tail. These distributions are shown in Figure 4.1. However, none of the Kolmogorov-Smirnov tests gave results that were consistent with the null hypothesis (that is, that the model is a statistically acceptable fit to all the spectra); even the variable persistent flux fits do not adequately fit the data, though they are an improvement on the old method. Some other

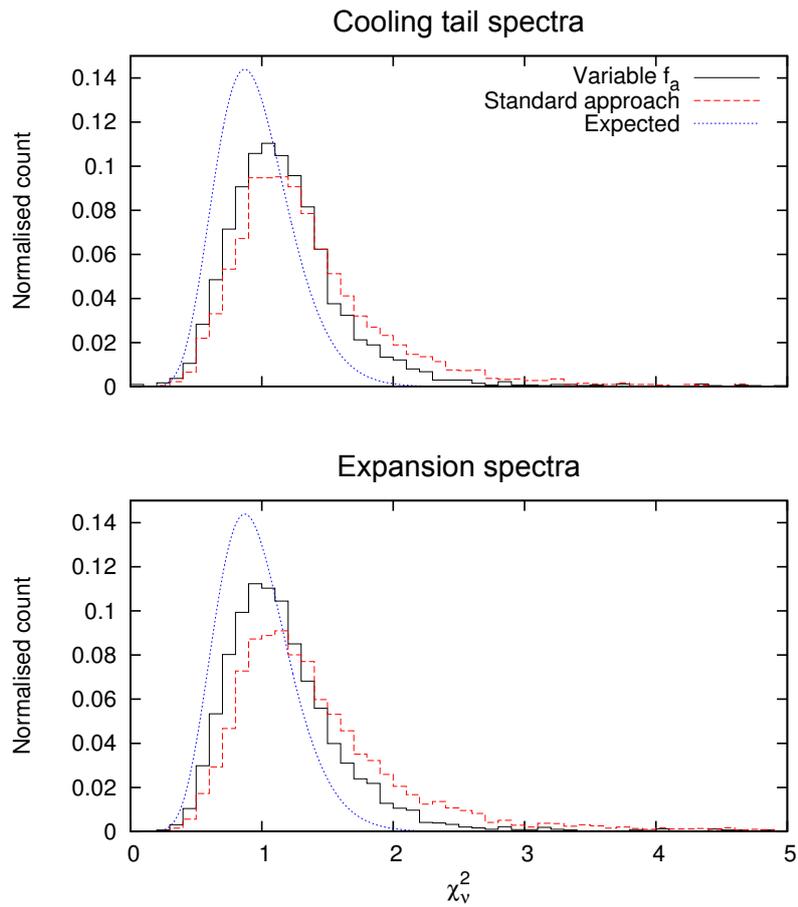


Figure 4.1 Distribution of reduced χ^2 for the variable persistent flux vs. standard approach spectral fits, for both cooling tail and Eddington-limited spectra. Also plotted is the theoretical distribution for 23 degrees of freedom, which is the mean for the standard approach fits (blue line). Allowing the persistent emission to vary significantly improves the fits, but there are still significant deviations from the expected distribution for a perfect model. The variable persistent flux and standard approaches are not consistent with being drawn from the same distribution using a Kolmogorov-Smirnov test.

systematic error must yet contribute to the discrepant spectral fits, perhaps deviations in the intrinsic burst spectrum from the assumed blackbody shape. Kolmogorov-Smirnov tests on the variable persistent flux method against the standard approach likewise indicates that the two methods are not consistent with each other, for either the cooling tail or Eddington-limited spectra.

I also investigated whether the degree of radius expansion correlates with the increase in persistent emission. I normalised the radius for an individual burst as the photospheric radius divided by the radius at the touchdown point, which corrects for the distance of the source, assuming the isotropy factor does not change. I compared the maximum normalised radius with the maximum f_a for each burst; see Figure 4.2. These two maxima do not necessarily occur at the same moment. I found only a slight relationship between the two quantities; a Kendall τ rank correlation test found $\tau = 0.138$ with 3.6σ (99.97%) significance. Though statistically significant, the small magnitude of the correlation suggests that obscuration of the inner parts of the disc does not greatly affect the observed persistent emission.

One possible contribution to the elevated persistent flux during contraction is the fallback of the extended atmosphere. During the expansion phase, the atmosphere can be driven off at about 10^{18} g s^{-1} (Weinberg et al. 2006). This is, naturally, around the order of \dot{M}_{Edd} , the Eddington accretion rate. As an order-of-magnitude estimate I make the assumption that half the expanded mass is driven off permanently. In a burst that is Eddington-limited for 5 seconds, the mass that is not expelled is about $2.5 \times 10^{18} \text{ g}$. There are approximately 40% more contraction spectra than expansion spectra. All the spectra at these high fluxes are taken at 0.25s intervals, so the atmosphere must generally contract more slowly than it expanded, by a factor of about 40%. It follows that the accretion rate due to the atmosphere falling back could be as much as $\dot{M}_{\text{Edd}}/3$. There are bursters for which the pre-burst accretion rate is only 1% of \dot{M}_{Edd} , so for these sources there should be a noticeable excess of contraction stage spectra with f_a around 20-30. I have not detected any such excess (see Figure 4.3), so I conclude that atmosphere fallback does not contribute significantly to the accretion luminosity. A careful comparison with non-PRE bursts would clarify this issue further, as by definition atmosphere fallback cannot occur in these events, and this will be investigated in Chapter 5.

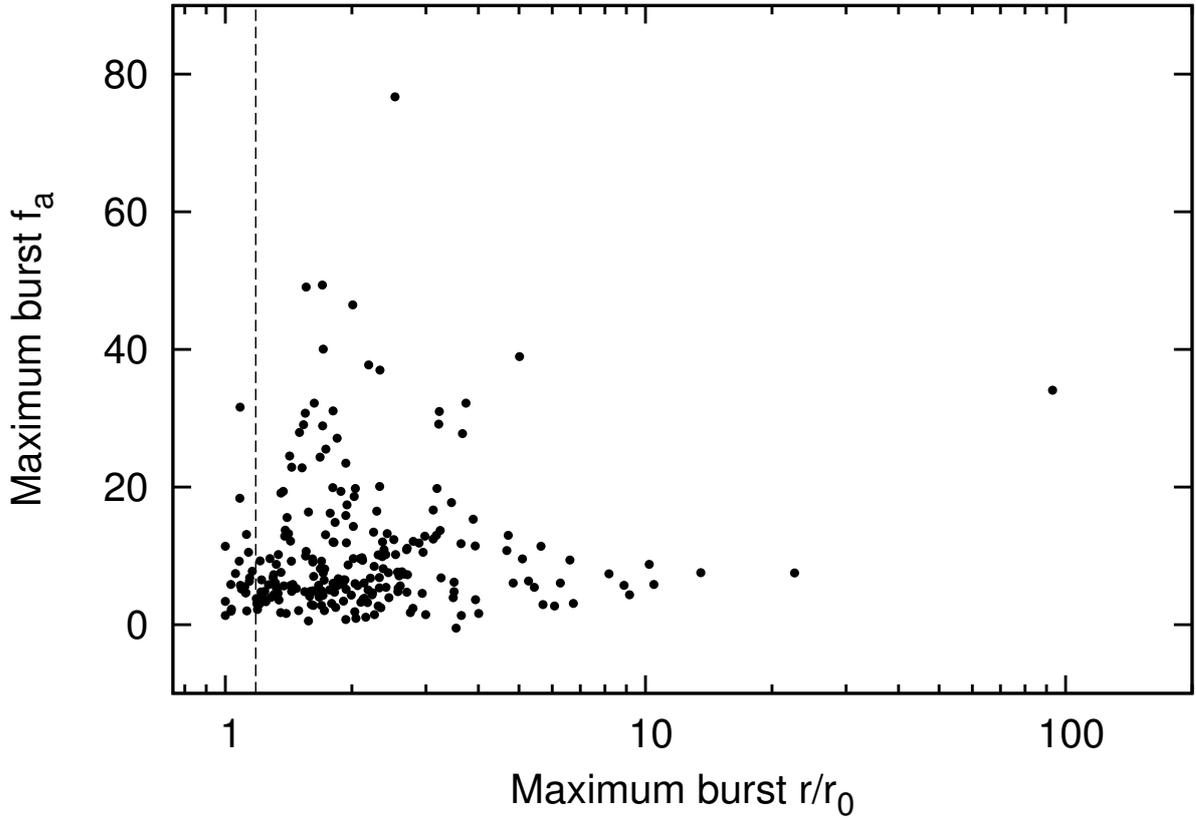


Figure 4.2 Maximum f_a plotted against normalised radius (photospheric radius divided by radius at touchdown) for each burst. The two maxima do not necessarily occur at the same time. The dashed line shows the expected location of an accretion disc boundary layer at $R = 1.18R_*$ (Popham & Sunyaev 2001). The majority of photospheric expansion events exceed this radius. There is no obvious relationship between maximum f_a and maximum expansion, nor is there a decrease in f_a outside the accretion disc radius, suggesting that obscuration of the disc by the expanding atmosphere does not significantly affect the observed accretion emission.

4.2 f_a as a function of burst stage

The distribution of persistent flux factor f_a for every burst spectrum (i.e., spectra from the Eddington-limited and cooling tail stages) is shown in Figure 4.3. The distribution has a wider spread for the Eddington-limited spectra, so that there is a larger fraction of high f_a spectra during radius expansion than for the cooling tail. However, the population of high f_a spectra in the cooling tail is not negligible. This suggests that elevated f_a is not only a result of the burst being Eddington-limited.

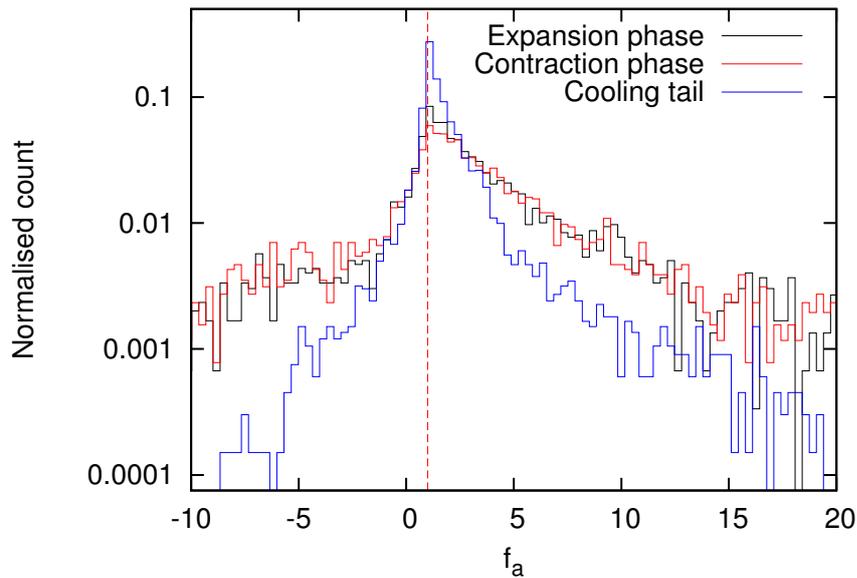


Figure 4.3 Distribution of measured pre-burst persistent emission factor f_a for cooling tail spectra and Eddington limited spectra for all PRE bursts. The cooling tail spectra are strongly peaked around 1, while the Eddington-limited spectra are not as strongly peaked. The distributions for the Eddington limited spectra have more prominent tails, indicating a large number of elevated f_a values, but there a significant number of high f_a values in the cooling tail also. The distribution for the expansion and cooling phases is consistent, suggesting that the fallback of the atmosphere does not contribute significantly to the emission.

There are 451 spectra with f_a less than zero, for which the f_a measurement is significant to more than 3σ (via the procedure detailed in Figure 2.5), and which occur in the Eddington-limited phase, out of 26,113. Of these, 333 spectra come from just six bursts. These six events are summarised in Table 4.1. By inspection of the flux-temperature curves I identified that four of the six are examples of *superexpansion* bursts (e.g., [in't Zand & Weinberg 2010](#)), particularly powerful PRE events for which the atmosphere is expanded to much larger radii than usual. In these events the temperature of the photosphere drops below about 0.2 keV, out of the detection band of *RXTE*, resulting in zero observed flux from the star. The other two are bursts with highly unusual flux-temperature curves and appear to consist of two consecutive expansions separated by a few seconds, followed by a cooling tail with larger blackbody radius than the maximum radius reached during expansion. These unusual bursts are shown in Figure 4.4.

Table 4.1. Bursts with many negative f_a burst spectra

Source	Date and time	G08 burst ID	Description
4U 1722–30	1996, Nov 8, 07:00	1	Super-expansion burst ^a
4U 2129+12	2000, Sep 22, 13:47	1	Super-expansion burst
2S 0918–549	2008, Feb 8, 03:02	5 ^b	Super-expansion burst ^c
4U 1722–30	2008, Mar 1, 16:18	4 ^b	Super-expansion burst ^a
XB 1832–330	1998, Nov 27, 05:45	1	Atypical burst profile
HETE J1900.1–2455	2010, Sep 20, 05:29	7 ^b	Atypical burst profile

^aThis burst was also studied by [in't Zand & Weinberg \(2010\)](#)

^bThis burst postdates G08; a consistent burst ID numbering scheme is assumed

^cThis burst was also studied by [in't Zand et al. \(2011\)](#)

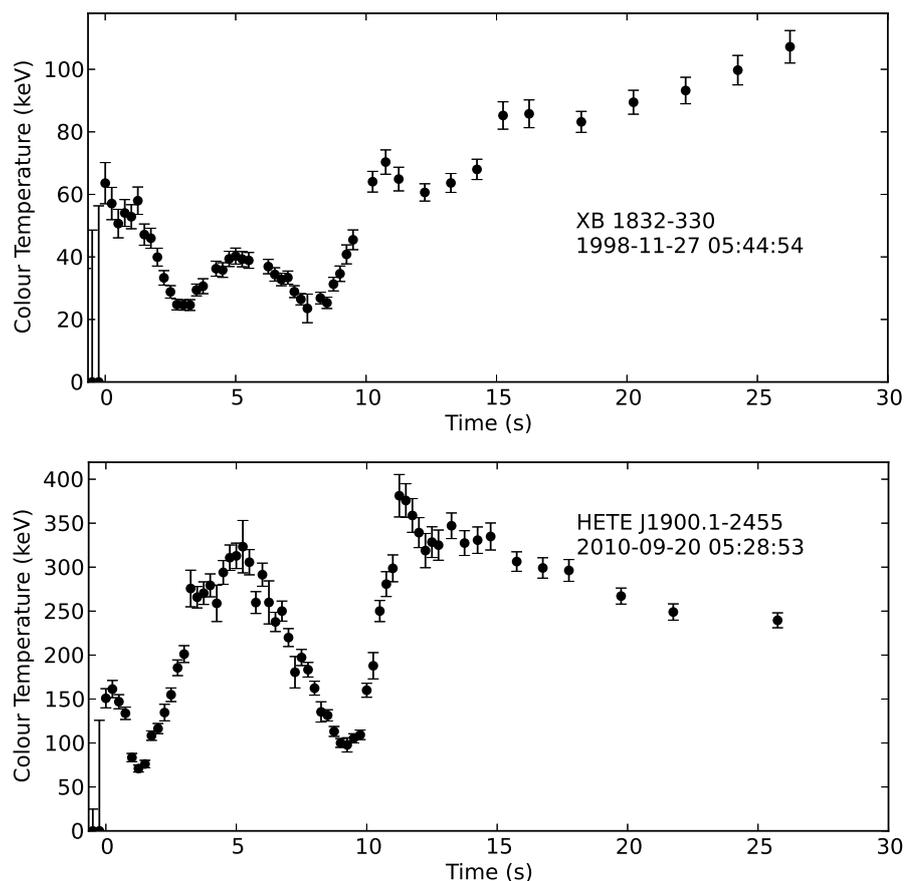


Figure 4.4 The two strange bursts with atypical profiles listed in Table 4.1. They appear to show two consecutive radius expansion periods separated by a few seconds, followed by a cooling tail with radius apparently larger than the maximum radius reached during the expansions.

4.3 Spectral shape changes

The new variable flux fitting model holds the shape of the persistent emission fixed, changing only its intensity. Since the persistent emission of LMXBs changes shape according to whether it is in a high or low state (Hasinger & van der Klis 1989), it is reasonable to think it may change shape temporarily if the accretion rate changes during a burst. It is possible that spectral shape changes in the persistent spectrum contribute to the high χ^2 in the approach. However, it is not clear that an accretion rate enhancement due to radiation drag will have the same effects on the persistent spectrum as the neutron star’s usual movement around the Z-track, since they have very different physical causes. Indeed, in’t Zand et al. (2013) found in their study of a burst from SAX J1808.4–3658 that the persistent spectrum became harder, whereas ordinarily they would expect increased accretion to soften the spectrum.

Deviations of the burst spectral component (that is, everything except persistent emission) from a blackbody could also be reflected by a changing f_a , especially if these deviations manifest as an excess of photons at high energy. Deviations from a pure black-body are frequently inferred during radius-expansion bursts (see, e.g., Kuulkers et al. 2002, 2003). Here I attempt to test the alternative hypothesis that the poor χ^2_ν values given by the standard approach arise instead from deviations of the burst component from a blackbody. The influence of these deviations on f_a measurements must be investigated, but this is difficult because the nature of the deviations is not known.

4.3.1 Phenomenological modification to burst spectrum

I take the approach of assuming the burst component retains a consistent spectral shape, which I model as a blackbody plus power law. I returned to the “prototype” burst (see §2.2) and fit the Eddington-limited and cooling tail spectra (stages 1, 2, and 3 in Table 2.2) with an absorbed blackbody plus powerlaw, with all parameters except N_H variable. The mean power law index of these fits was 1.93 ± 0.68 . I then fit the spectra again, this time holding the power law index fixed at this value but allowing the normalization to still vary. As I am here assuming that the power law is intrinsic to the burst emission, rather than being a separately varying component, I must specify the power law normalization so

that it contributes a fixed proportion of the total flux. Using a linear fit, I found that the normalization of the power law was best described by

$$K_{PO} \approx \frac{K_{bb} \times kT^4}{3212}, \quad (4.3.1)$$

where kT is the blackbody temperature in keV and K_{bb} is the normalization of the blackbody in $\text{km}^2/100\text{kpc}^2$. I fit the spectra a third time, this time tying the power law normalization to the blackbody parameters with the relation given by equation 4.3.1, which fixes the shape of the burst emission.

I found the mean χ^2_ν was 1.27 ± 0.43 for the spectra of the Jun 15, 2001 burst from 4U 1636–536 using this fitting method, compared with 1.69 ± 0.58 and 1.21 ± 0.37 for the standard and variable persistent methods respectively. Thus, for this burst, the phenomenological model plus fixed persistent component fits the data nearly as well as the variable persistent model. I then used the same model, with the same power law index and normalization relation, to fit every stage 1, 2, and 3 spectrum from all radius expansion bursts in the catalog. Restricting these to spectra where XSPEC did not encounter fitting errors for any of the three methods, I found that the phenomenological model gave a mean χ^2_ν of 1.43 ± 1.10 , compared with 1.45 ± 0.82 and 1.21 ± 0.60 for the standard and variable persistent methods respectively. The number of degrees of freedom are typically in the range 20 to 25. The phenomenological model, plus fixed persistent emission, is thus globally no better than the standard approach. If there is a non-blackbody contribution to the burst component, it must differ between sources and change from burst to burst. Furthermore, it would also have to resemble the persistent emission for every burst, at least superficially, or the variable persistent flux method would not be able to consistently improve the reduced χ^2 ; and this itself hints at a relationship with the persistent emission. Clearly any further investigation into such deviations must be physically motivated rather than phenomenological.

4.3.2 Does a modification to the burst spectrum remove the need for variable persistent emission?

I then tested whether introducing a phenomenological change to the burst spectrum removes the need for a variable persistent emission. I fit the spectra of the Jun 15, 2001 burst from 4U 1636–536 using the variable persistent model, but replacing the blackbody burst component with the phenomenological spectral model. If the initial detection of a variable persistent flux merely reflected the non-blackbody character of the burst component, then I would expect the detection to disappear. I still detect f_a to vary significantly, though the values are consistently lower by roughly 20% than for the original fits. Furthermore, the mean χ^2_ν for these fits was 1.17 ± 0.35 , suggesting that allowing the persistent emission to vary improves the fits even when a non-blackbody model for the burst component is assumed. Again, the number of degrees of freedom here are generally in the range 20–25. These results suggest very strongly that the detection of f_a cannot be attributed to a confounding spectral effect intrinsic to the burst component.

4.3.3 Correlation with burst flux?

If the $f_a > 1$ results are arising from a non-blackbody contribution to the burst spectrum then I should expect f_a to increase with burst flux. Suppose that a non-Planckian part of the burst component contributes a fixed fraction of the burst component flux and that the variable persistent flux fits are trying to remove it. Then if the burst flux doubles, I would require f_a to double also in order to fit out the non-Planckian part. I investigated this by comparing f_a against F/F_{Edd} , where F_{Edd} is the source’s Eddington flux (see Appendix A) for all Eddington-limited and cooling tail spectra. It is clear from Figure 4.5 that there is no relationship between the quantities. To test whether plotting all the points together obscures a relationship present in individual bursts I selected the Eddington-limited and cooling tail spectra from the prototype burst and three randomly selected other bursts, and performed Kendall rank correlation tests (Kendall 1948) on their f_a against their blackbody fluxes.

No statistically significant correlation was present in any of these bursts. If the hypothesised non-blackbody contribution is present at the high energy tail, then I would expect an

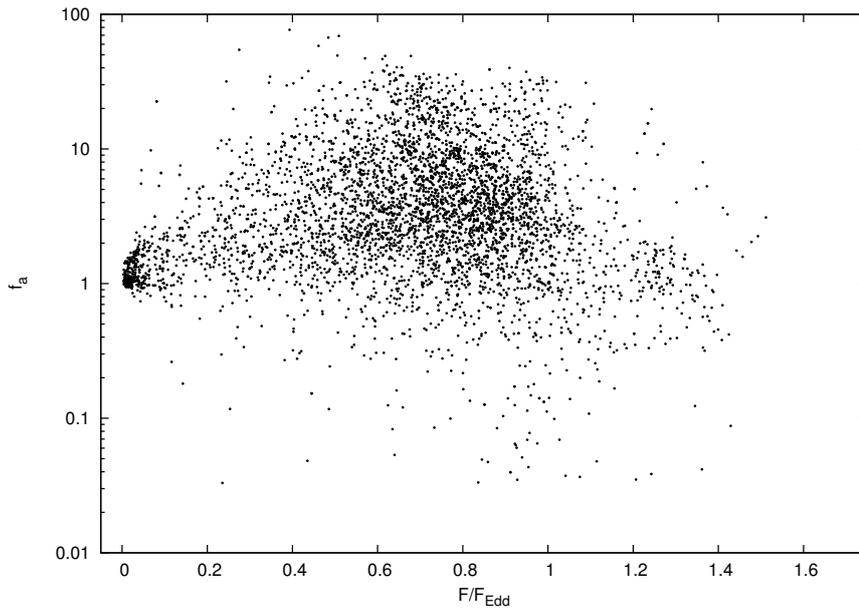


Figure 4.5 f_a against normalized burst flux for all Eddington-limited and cooling tail spectra. There is no obvious correlation between the two quantities. There is a clearly visible clump at low flux and $f_a \approx 1$, showing that f_a returns to its original level after the burst. There are also many points with high f_a and normalized fluxes between 0.2 and unity, indicating that f_a can remain elevated a considerable way down the cooling tail.

inverse correlation between blackbody temperature and f_a . At low temperatures the blackbody component of the overall spectrum drops out of the band detectable by *RXTE*, leaving only the hard tail to be fit. However, I find a slight *positive* correlation of $\tau = 0.031$ at 4.2σ significance using a Kendall rank correlation test.

4.3.4 Neutron star atmosphere spectra

Finally, I further tested whether enhanced f_a is attributable to the burst component of the spectrum deviating from a blackbody by comparing observed spectra to theoretically motivated ones. [Suleimanov et al. \(2011\)](#) calculated spectra for neutron star atmospheres, with a range of chemical compositions and surface gravities. These spectra are different from blackbodies. I tested whether the differences are detectable in *RXTE* data with the standard technique of simulating model spectra, fitting them with a blackbody, and testing whether any differences manifest as unacceptably big χ^2 . These models have stellar radius, distance and burst luminosity as their variable parameters, and composition and surface gravity ($\log g \in [14.0, 14.3, 14.6]$) as discrete parameters.

We took each possibility for the surface gravity, and either solar or pure helium composition, and took the distance to be 8 kpc. For each combination I generated a two-dimensional grid in stellar radius and luminosity and for each grid point I generated a simulated spectrum of the corresponding atmosphere. I then fit these with the usual blackbody model. I did not find any significantly poor χ^2 values, indicating that these spectra are indistinguishable from blackbodies.

4.4 Comparison with theory

W92 performed one-dimensional computer simulations of radiation drag on an accretion disc during a Type I burst. He found that accretion enhancements of up to 100 times the pre-burst levels could be achieved. I have compared the f_a measurements to his results to test whether interpreting increased persistent flux as increased accretion rate is reasonable.

Table 1 of W92 lists peak accretion enhancements for the computed models; in Figure 4.6 I compare the maximum f_a measured for each burst against γ , the pre-burst accretion flux as a fraction of the Eddington flux of the burst source. The Eddington flux for each neutron star is the mean peak flux of every PRE burst observed from that star; see Appendix A for the values and details of their calculation. Following G08 I measure γ by integrating the chosen pre-burst persistent model between 2.5 and 25 keV and dividing the resulting flux by the burst source's Eddington flux. W92's Table 1 lists related quantities. W92's models begin with a non-rotating neutron star with radius 9km, γ of roughly 0.3, and disc viscosity parameter, β , of 10^{-4} . They then allow the spin frequency, accretion rate, β , and radius to vary in turn, while holding the other quantities fixed at their original values.

The disc viscosity parameter β (Coroniti 1981) is similar to the Shakura-Sunyaev disc viscosity (Shakura & Sunyaev 1973) but relates the viscosity to the gas pressure rather than the total pressure, which differs in radiation-pressure dominated discs. This β is not to be confused with the radiation drag parameter β described in Chapter 5. W92 gives accretion rates in dimensionless units: $\dot{m} = \dot{M}c^2/L_{\text{Edd}}$, where \dot{M} is the mass accretion rate and L_{Edd} is the Eddington luminosity, whereas I give accretion rates in terms of the energy

release. Since the mass M_* and radius R_* of the neutron stars in W92 are specified, I have

$$\gamma = \frac{GM_*}{c^2 R_*} \dot{m}. \quad (4.4.1)$$

If the quantity f_a is assumed to be entirely an accretion enhancement, then it corresponds identically to their quantity $\Delta\dot{M}_*/\dot{M}_*$ and is predicted to be inversely correlated with pre-burst accretion rate, accretion disc viscosity parameter β , neutron star spin frequency, and neutron star radius. All of W92's models assume a neutron star of mass $1.4M_\odot$.

The results in Figure 4.6 show a decrease in peak f_a with increasing γ , and the slope is consistent with the predictions of W92's three points representing a nonrotating star with increasing accretion rate, while the peak f_a values are significantly lower than those predicted for a nonrotating neutron star. If I assume a moderate β and rotation frequency of about 300 to 600 Hz (e.g., [Muno et al. 2001](#), [Watts 2012](#); see also Appendix A), then not only the observed correlation of peak f_a with γ , but also the normalization, appear to be consistent with W92's predictions (see Figure 4.6). The models predict that stellar rotation period and disc viscosity parameter β have a large influence on peak f_a , but that neutron star radius apparently has little influence. The upper edge of the measured points appears to be very roughly consistent with $\gamma f_a \lesssim 1$, implying that the accretion luminosity cannot greatly exceed L_{Edd} . This is consistent with the predictions of [Burger & Katz \(1983\)](#) and [Miller & Lamb \(1996\)](#), who find that \dot{M}_{Edd} is a natural upper limit on the accretion rate.

In Figure 4.7 I plot maximum f_a against spin rate for all neutron stars in the catalog whose spin rate is known. These are listed in Appendix A. W92's models predict a gradual decrease of maximum f_a with spin frequency. The data shows maximum f_a values consistent in magnitude to W92's models, but it is difficult to discern any trend, because of the large scatter in individual sources (due to the additional dependence on γ and the disc viscosity) and the small range of known neutron star spin frequencies in PRE bursting systems. I would require PRE bursts from neutron stars rotating at less than 100 Hz to better constrain this relationship. A further difficulty is that W92's models do not reach the Eddington limit, so I am comparing PRE bursts in observations to non-PRE bursts in models.

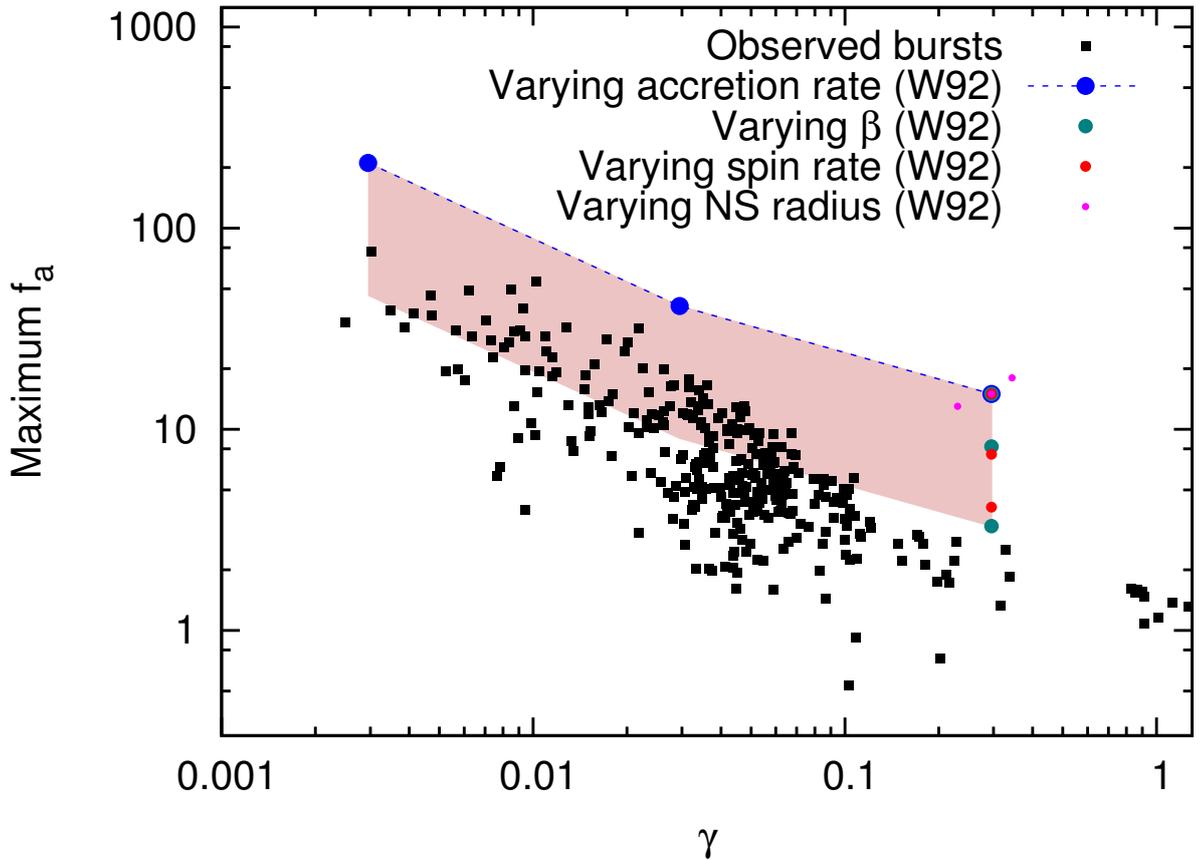


Figure 4.6 Maximum f_a against γ (the pre-burst accretion rate as a fraction of \dot{M}_{Edd}) for all PRE bursts (black squares). Also plotted are the results of ten computer simulations performed by W92 (see their Table 1). The observed bursts show generally lower f_a , but the same slope, as W92's nonrotating neutron star model. If a moderate β and spin frequency of 300 to 600 Hz is assumed, then the results of this analysis appear to agree with W92. The shaded region indicates the approximate area spanned by W92's models. The bottom edge of the shaded box is uncertain due to the fact that none of W92's models vary both rotation rate and disc viscosity simultaneously, but it may extend further down than shown here.

As the W92 models are one-dimensional, care must be taken in applying them to accreting neutron star systems in which three-dimensional effects are likely to be important. Further theoretical studies, preferably in three dimensions, would be very valuable. I attempt to address this shortcoming in Chapters 5 and 6.

4.5 Luminosity of bursters in globular clusters

Radius expansion bursts from sources in globular clusters provide a direct measure of the Eddington *luminosity* of the star, as opposed to just its Eddington *flux*, because the

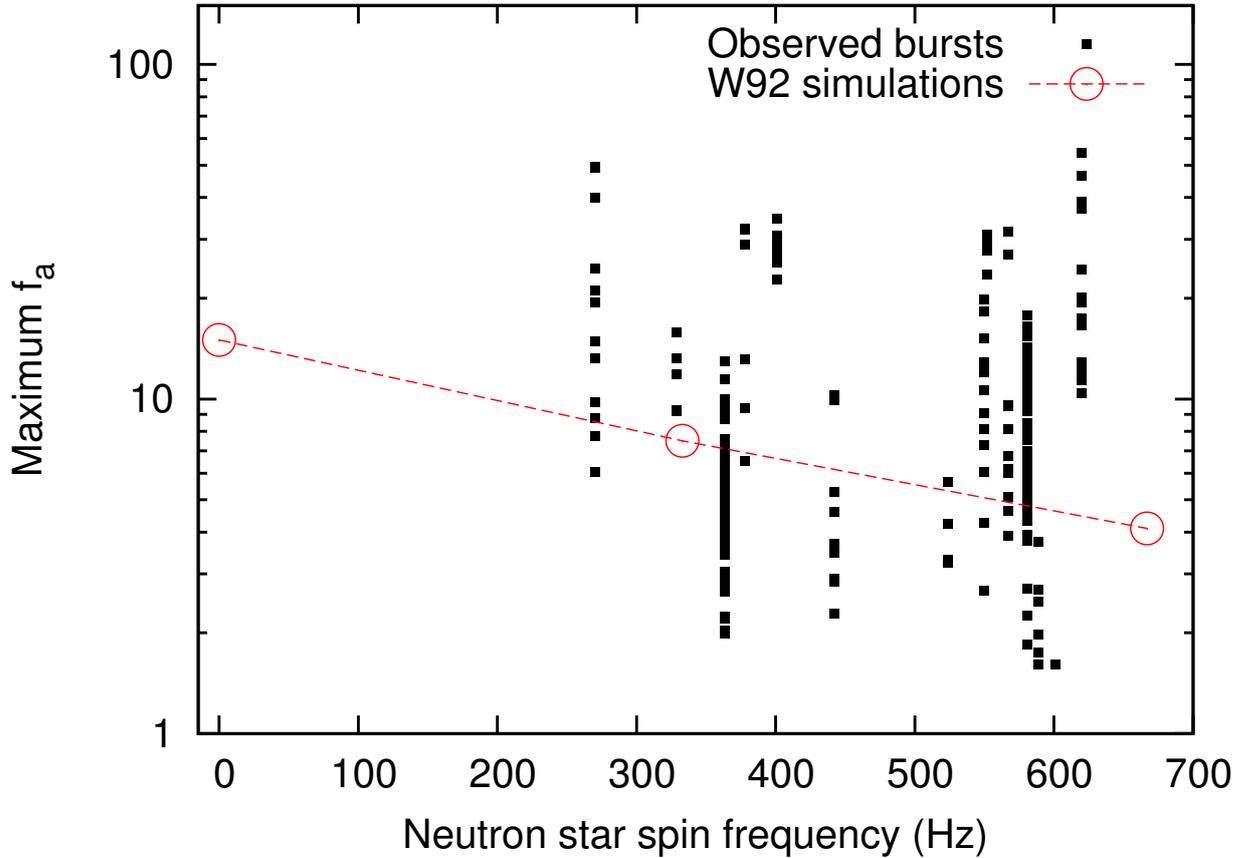


Figure 4.7 Maximum f_a against neutron star spin frequency for all PRE bursts arising from sources whose spin frequency is known (see Appendix A). Also plotted are the results of three computer simulations performed by W92 with spin frequencies of 10^{-5} , 333, and 667 Hz. Although there is large scatter due to the dependence of f_a on pre-burst \dot{m} and β , these results are consistent in magnitude with W92's predictions. The anti-correlation of peak f_a with spin period predicted by theory is not obviously visible in the data.

distances to the globular clusters can be independently measured. In a study of 46 radius expansion bursts from sources in twelve globular clusters (Kuulkers et al. 2003) found that two-thirds of sources reach the *same* peak luminosity of $37.9 \pm 0.15 \times 10^{37}$ erg s $^{-1}$. The remaining sources are thought not to have reached the same critical luminosity because they are high inclination systems, so that we do not observe all the flux, or because they were bursting in a hydrogen rich atmosphere, so that their Eddington luminosities is 70% lower than for hydrogen-poor material. I have repeated this study using radius expansion bursts from globular cluster sources detected by *RXTE* and *BeppoSAX*.

Table 4.2. Globular cluster burst sources and their distances

Burst source	Globular cluster	Distance (kpc)	Ref.
4U 0513–40	NGC 1581	12.1 ± 0.3	1
4U 1722–30	Terzan 2	$9.5^{+2.5}_{-2.0}$	1
EXO 1745–248	Terzan 5	$8.7^{+3.3}_{-2.4}$	1, 2
4U 1746–37	NGC 6441	$11.0^{+0.9}_{-0.8}$	1, 3
GRS 1747–312	Terzan 6	$9.5^{+3.3}_{-2.5}$	1
SAX J1748.9–2021	NGC 6440	$8.4^{+1.5}_{-1.3}$	1
4U 1820–30	NGC 6624	7.6 ± 0.4	4
XB 1832–330	NGC 6652	9.6 ± 0.4	1, 5
4U 2129+12	NGC 7078	10.3 ± 0.4	1

References. — 1. Harris (1996); 2. Cohn et al. (2002); 3. Pritzl et al. (2001); 4. Heasley et al. (2000); 5. Chaboyer et al. (2000)

4.5.1 Data

I have studied 106 radius expansion bursts from 9 sources. All are in different globular clusters, as indicated in Table 4.2. The distances to these globular clusters are taken from Kuulkers et al. (2003) and references therein. The peak fluxes, as measured using the standard approach, F_{peak} of the radius expansion bursts and their uncertainties are taken from the MINBAR catalogue. Many of the bursts studied in this section were also studied in Kuulkers et al. (2003), but the details of instrument calibration and spectral analysis are different.

I calculated the peak luminosities of the radius expansion bursts according to

$$L_{\text{peak}} = 4\pi d^2 F_{\text{peak}}, \quad (4.5.1)$$

where d is the distance to the source. Errors on L_{peak} are propagated from errors on d and F_{peak} in the usual way:

$$\Delta^{\pm} L_{\text{peak}} = L_{\text{peak}} \sqrt{\left(\frac{\Delta F_{\text{peak}}}{F_{\text{peak}}}\right)^2 + 2 \left(\frac{\Delta^{\pm} d}{d}\right)^2}, \quad (4.5.2)$$

where the superscripted \pm accounts for the different upper and lower uncertainties for the distances of some of the globular clusters.

4.5.2 Statistical method

I follow the approach suggested in [Press \(1996\)](#) for combining measurements of a quantity that are potentially inconsistent. The probability that a given burster will have an Eddington luminosity of L_{Edd} , given the burst data that we have for it, is given by a proportionality:

$$P(L_{\text{Edd}}|D) \propto P(L_{\text{Edd}}) \sum_p P(p) \prod_i [pP_{Gi} + (1-p)P_{Bi}], \quad (4.5.3)$$

where $P(L_{\text{Edd}})$ is our prior on the Eddington luminosity of the source and $P(p)$ is the prior on the probability that a given measurement will have appropriate error bars. These are discussed in more detail in [§4.5.3](#). P_{Gi} and P_{Bi} are the probability distributions of good and bad measurements on the i th measurement; they are given by

$$P_{Gi} = \exp \left[-\frac{(L_{\text{Edd},i} - L_{\text{Edd}})^2}{2\sigma_i} \right], \quad (4.5.4)$$

and

$$P_{Bi} = \exp \left[-\frac{(L_{\text{Edd},i} - L_{\text{Edd}})^2}{2S_i} \right]. \quad (4.5.5)$$

The quantities $L_{\text{Edd},i}$ and σ_i are the Eddington luminosity of the i th burst and its uncertainty, and S_i is a large, but finite, number describing a plausible range in which a wrong result could have been published without being immediately rejected as obviously wrong. Although S_i is essentially arbitrary the final result is somewhat sensitive to its value. I assumed $S_i = 6\sigma_i$.

I performed this procedure for every source individually, and took L_{Edd} to be the location of the peak of the probability distribution. Since some of the probability distributions are significantly skewed, I have estimated errors on L_{Edd} as follows. I first calculated the areas of the probability distribution above and below the peak. Then I calculated the point where 68.26% of that area was contained between the point and L_{Edd} .

As is customary in Bayesian probability calculations, I did all the computations as proportionalities and left the normalisations to the end.

4.5.3 Priors

The prior on the Eddington luminosity of each source is not well known. I can safely assert that Eddington luminosities are greater than zero, and I assume all other possibilities are equally likely.

$$P(L_{\text{Edd}}) \propto \begin{cases} 1, & 0 < L_{\text{Edd}} \\ 0, & \text{otherwise} \end{cases} \quad (4.5.6)$$

The prior on p , described in [Press \(1996\)](#) as the community-wide probability of publishing results with error-bars appropriate to the measurement, is also not well known. That is, $1 - p$ is the probability that the measurement is wrong, including the possibility that the measurement is affected by systematic errors that are not negligible compared with the published uncertainties. The result is somewhat sensitive to $P(p)$ so I must give some thought to how I deal with it. I assume that there is no chance that the community is *incapable* of ever obtaining appropriate error bars and, similarly, no chance that it is completely infallible; thus, $P(0) = P(1) = 0$. I have chosen a prior that peaks at $p = 0.75$, see also [4.8](#):

$$P(p) \propto \begin{cases} p^2, & p < 0.75 \\ 9p^2 - 18p + 9, & 0.75 \leq p \end{cases} \quad (4.5.7)$$

4.5.4 Results

Probability distributions for the Eddington luminosities for individual sources are given in [Figure 4.9](#). It is clear that they are not compatible with a single value.

The effect of choosing $S_i = 6\sigma_i$ can be most clearly seen in the panel for XB 1832-330. There is a sharp central Gaussian-like peak representing the possibility that the single data point is correct and has appropriate error bars, as one would expect. The probability distribution has prominent tails- these represent the possibility that the true L_{Edd} is, say, $60 \times 10^{37} \text{ erg s}^{-1}$ and that the single data point is erroneous with wrong error bars, but still

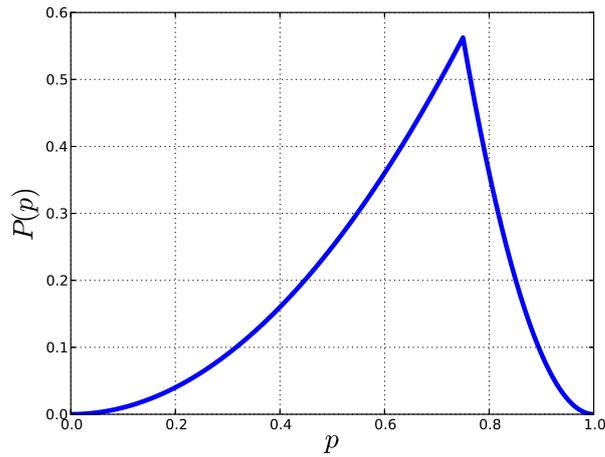


Figure 4.8 The prior $P(p)$, the prior belief that the community has a probability of p of publishing a given measurement with appropriate error bars. It has the property that $P(0) = P(1) = 0$, i.e., the community is neither infallible nor always wrong, and I have set the prior to peak at $p = 0.75$. See also Equation 4.5.7.

near enough to the true value that the result was not immediately rejected as obviously erroneous.

In Figure 4.10 I show the Eddington luminosities for all nine sources against the theoretical Eddington luminosity of a neutron star with mass $1.4M_{\odot}$ and radius 10 km, for hydrogen-rich and hydrogen-poor atmospheric compositions. It is clear from this Figure that the Eddington luminosities are not consistent with the predicted Eddington luminosity of the “canonical” neutron star. It may be possible to explain the low Eddington luminosity sources, such as 4U 1746–37, as a result of hydrogen-rich atmospheric composition or because they are “dippers”, that is, systems viewed almost edge-on, so that the accretion disc obscures much of the flux. It is clear that the Eddington luminosities for the remaining sources are systematically higher than the theoretical prediction for the “canonical” neutron star of $M = 1.4M_{\odot}$ and $R = 10$ km, with a hydrogen-poor atmosphere. Perhaps the neutron stars are more massive than $1.4M_{\odot}$; if so, then their calculated theoretical Eddington luminosities would better match observation, as shown in Figure 4.11.

The width of the Eddington luminosity bands in Figures 4.10 and 4.11 reflect the difference in Eddington luminosity for a PRE burst that barely leaves the surface ($R_{\text{photosphere}} \approx R_{*}$) compared to one that reaches a large distance from the star ($R_{\text{photosphere}} \gg R_{*}$), due to gravitational redshift as explained in Equation 1.3.1. Increasing the mass of the star moves

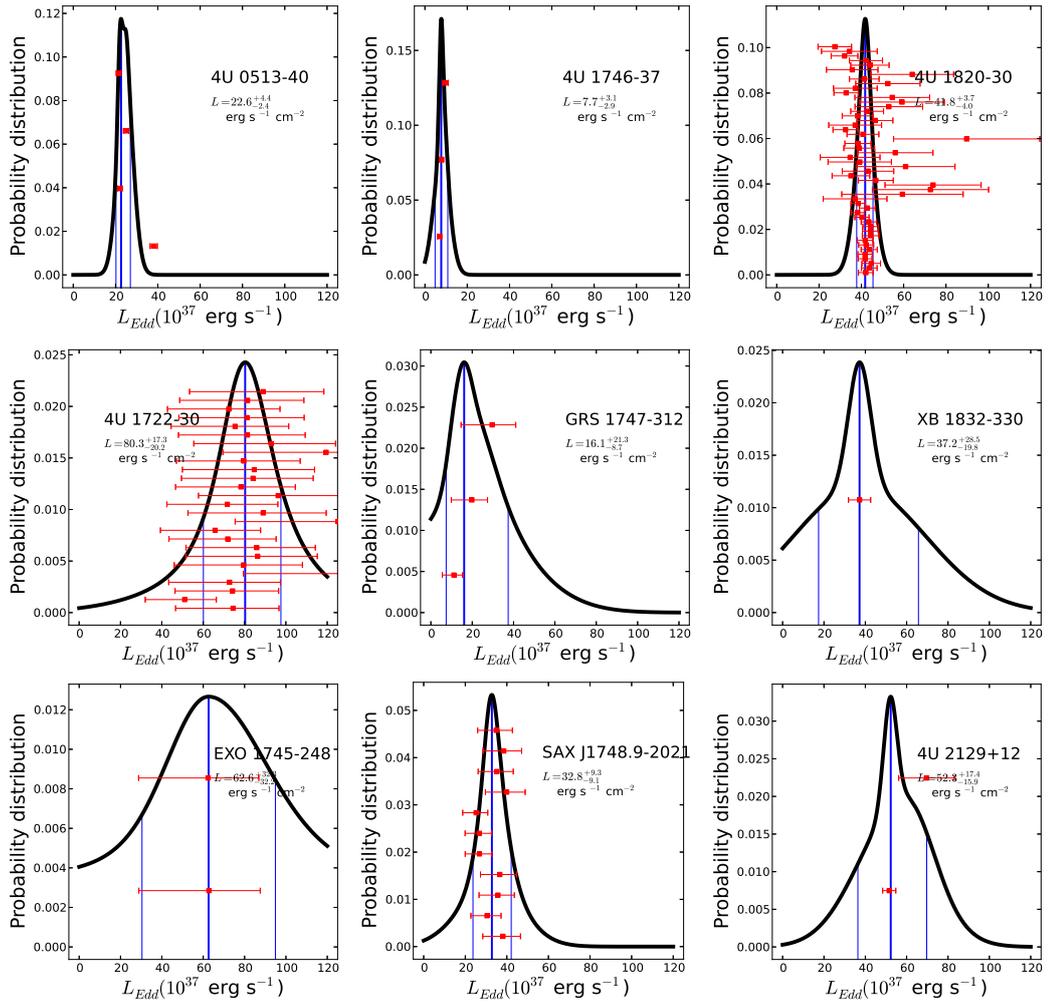


Figure 4.9 Probability distributions of L_{Edd} for the nine globular cluster sources (black lines), with peak probabilities and 1σ intervals indicated as blue lines. The red points with error bars are the individual peak luminosity measurements, as calculated in Equations 4.5.1 and 4.5.2. It is clear that they are not compatible with a single value.

the bands upwards and widens them. Decreasing the neutron star radius moves the lower edges of the bands downwards without affecting the upper edges.

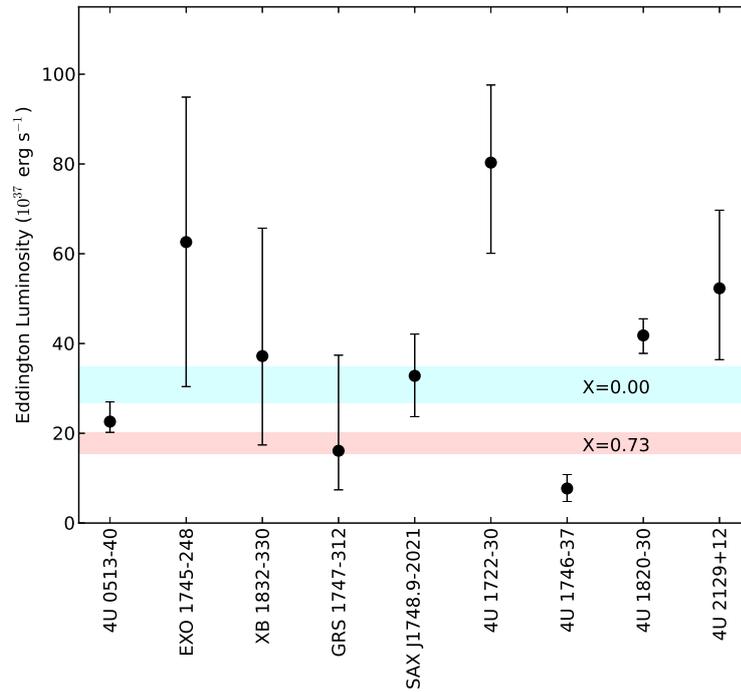


Figure 4.10 Eddington luminosities for the globular cluster sources against the theoretical Eddington luminosity of a star with mass $1.4M_{\odot}$ and radius 10 km, for hydrogen-rich (red band) and hydrogen-poor (blue band) atmospheric compositions. The Eddington luminosities are not consistent with a single value, even if dippers and hydrogen-rich sources are accounted for.

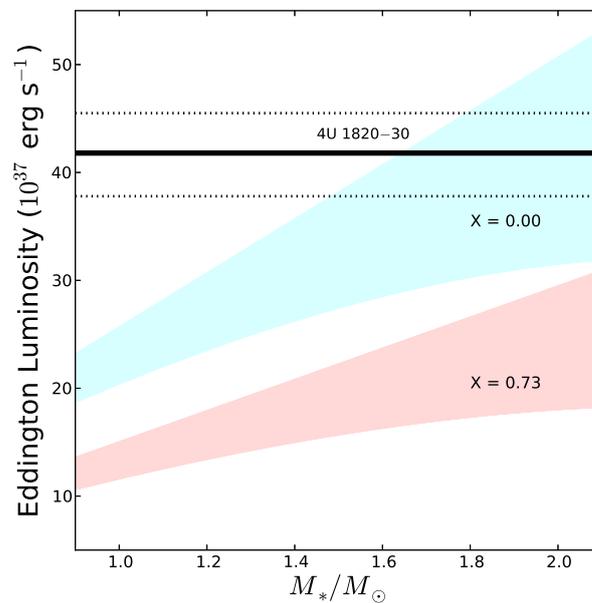


Figure 4.11 The effect of neutron star mass on the theoretical Eddington luminosity, showing the increase in Eddington luminosity with increasing neutron star mass. As in Figure 4.10 the red and blue bands represent hydrogen-rich and hydrogen-poor atmospheric composition. The upper and lower edges of the bands represent very vigorous and very feeble radius expansion respectively. Also plotted is the L_{Edd} for 4U 1820–30; if the mass of this star were above about $1.6M_{\odot}$ its observed L_{Edd} would be consistent with theory.

Chapter 5

Observational Results: Non-Eddington Limited Bursts

-The same procedure as last year, Miss Sophie?

-The same procedure as *every* year, James!

James and Miss Sophie,

Dinner for One

Photospheric radius expansion bursts change the structure of the neutron star photosphere, and this change is likely to cause the spectrum of the photosphere to deviate from its usual near-blackbody shape. Such changes are unrelated to the accretion spectrum. As pointed out in Chapters 2 and 3, these may confound spectral analyses of persistent emission during a burst. Such complications can potentially be avoided by considering non-PRE bursts, but care must be taken to account for lower sensitivity due to their intrinsically lower flux.

5.1 Data analysis

I initially considered the full sample of 2,083 type I bursts observed by *RXTE*. A total of 168 bursts were excluded as being source confused. I also excluded 35 bursts whose radius expansion status was unclassifiable due to data gaps or very low flux, and 58 bursts classified as “marginal” (i.e. satisfying only some of the criteria of §2.3 in G08).

My catalog contains an additional 16 type I bursts from the Rapid Burster taken during offset pointings to avoid confusion with the nearby 4U 1728–34. These were identified as type I bursts by visual inspection of their light curves (Bagnoli et al. 2014).

For each burst I fit the persistent emission with a set of nine alternative models in turn. These are summarised in Table 5.1. The first two of these include accretion rate explicitly as a variable parameter, making it possible to investigate the relationship between the normalisation factor f_a and \dot{M} . I then selected the fit that gave the best (i.e. lowest) χ^2_ν to represent the persistent spectrum for that burst. As shown in Figure 5.1, this suite of models provide good fits, and a Kolmogorov-Smirnov test gave a 97% probability that the two curves are consistent with having been drawn from the same distribution ($D = 0.11$) so I regard these persistent models acceptable in the analysis of the bursts themselves.

I assigned each model a distinct numerical label and performed a K-S test on the resulting distributions, with $D = 0.03$ and a 95% probability that both are consistent with having been drawn from the same distribution. This result shows that there is no difference in the distributions of best-fitting persistent models between the PRE and non-PRE bursts.

I note that a gas pressure dominated accretion disc (`wabs*diskm`) is the preferred model for no persistent emission spectrum, indicating that this is not a good description of accretion discs in LMXB systems. Instead, I expect the disc to be radiation dominated, and possibly showing hard emission.

Profiles of fit parameters, using both methods, for three example bursts are shown in Figure 5.2. I have selected bursts from three different sources to illustrate the results. These plots clearly show that f_a is enhanced to several times the pre-burst level, as found for radius expansion bursts in Chapter 3. This result is typical for all the bursts in the analysis. The f_a for spectra preceding the burst appear to be slightly lower than unity, but this is an artefact of fitting these spectra with a model that includes a burst component which in reality is absent before the onset of nuclear burning. This causes some of the persistent flux to be misidentified as burst emission. Fitting the pre-burst spectra with just a normalisation-variable persistent model gives results consistent with f_a equal to unity.

In Figure 5.3 I show the distributions of χ^2_ν for the variable persistent normalisation fits and the standard approach fits, compared with a theoretical distribution of χ^2_ν for a model that

Table 5.1. XSPEC models for fitting the persistent emission

XSPEC model	Number of spectra	Notes
wabs*disko ^a	160	
wabs*(bbodyrad+powerlaw)	425	nH fixed to literature values
wabs*(bbodyrad+powerlaw+gaussian)	271	nH fixed, Gaussian energy set to 6.4keV
wabs*(compTT) ^b	39	nH allowed to vary
wabs*(bbodyrad+powerlaw)	46	nH allowed to vary
wabs*(gauss+bremss)	321	all parameters variable ^c
wabs*(bbodyrad+diskbb) ^d	183	all parameters variable ^c
wabs*(compTT+gaussian) ^e	314	nH fixed, Gaussian energy set to 6.4keV
Total usable bursts	1759	
Rejected due to source confusion	168	Other active sources in field
Background data missing or unusable	57	
No good persistent model fit	6	Minimum $\chi^2_\nu > 3.5$
Marginal radius expansion	58	
Radius expansion status unclassified	35	
Total unusable bursts	324	

^aSee [Stella & Rosner \(1984\)](#)

^bSee [Titarchuk \(1994a\)](#)

^cAll parameters are variable for the persistent emission fits. Their values are subsequently frozen for the burst spectral fits in this chapter

^dSee XSPEC manual (<http://heasarc.gsfc.nasa.gov/xanadu/xspec/manual/manual.html>) and references therein

^eSee [Peille et al. \(2014\)](#)

adequately describes the data. Fits from both PRE bursts and non-PRE bursts are shown. I performed Kolmogorov-Smirnov tests on the measured distributions of χ^2_ν against the theoretical. D and p values are listed on Figure 5.3, with the variable fits listed first. These results confirm what is visually evident in the Figure: the variable persistent normalisation fits significantly improve the quality of the spectral fits for both radius expansion and non radius-expansion bursts. The spectral fits for PRE bursts are generally poorer than non-PRE bursts for both fitting methods and, although the variable persistent normalisation method improves the distribution of χ^2_ν , the deviation from a model that is statistically consistent with the data is still present. It is clear that radius expansion introduces a significant spectral effect on top of the variations induced by enhanced persistent emission.

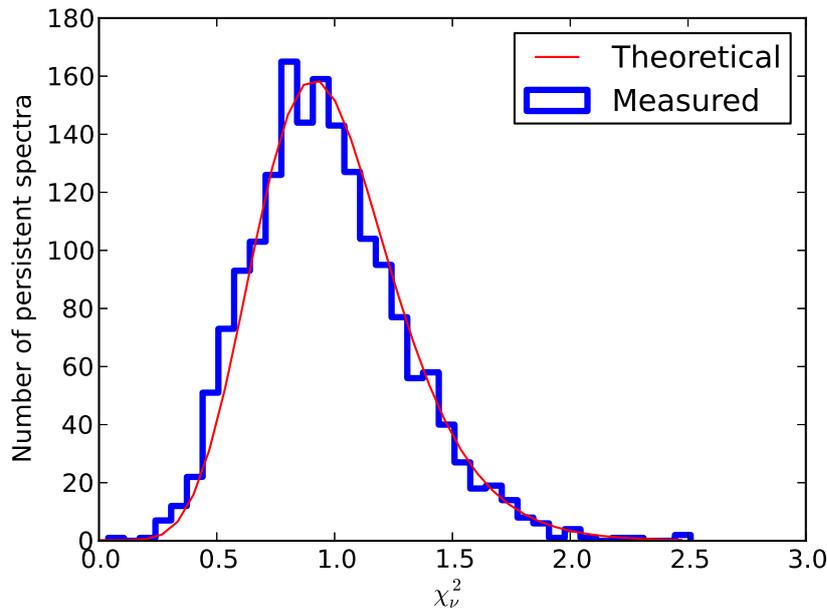


Figure 5.1 Distribution of χ^2_ν for the spectral fits to the persistent emission (histogram) compared to a theoretical distribution of χ^2_ν for a collection of spectra with the same number of degrees of freedom (curve). A K-S test gave a 97% probability ($D = 0.11$) that the two curves are consistent with having been drawn from the same distribution, indicating that the suite of models for the persistent emission spectra is adequate for use in subsequent work.

The second and third panels show that cooling tail spectra of PRE bursts are more poorly fit than spectra from non-PRE bursts. This may imply that PRE affects the disc and/or photosphere in a manner that persists for some time after the photosphere has returned to the surface of the star.

As in Chapter 4, there are generally 20-25 degrees of freedom in the spectral fits.

5.1.1 Tests of significance

I quantified the relative improvement of the variable persistent flux treatment over the standard analysis by estimating the Bayes factor, as follows. I assume that the probabilities calculated from the K-S statistics represent the global likelihood for each model. That is, I take the probability to be $p(D|M_i, I)$, where D is the set of χ^2 values from the model spectral fits, M_i represents the i th fitting approach (model), and I the priors.

Perhaps this interpretation is not strictly correct, but even so one would expect that the K-S probabilities are proportional to the actual likelihoods, so that a ratio of probabilities

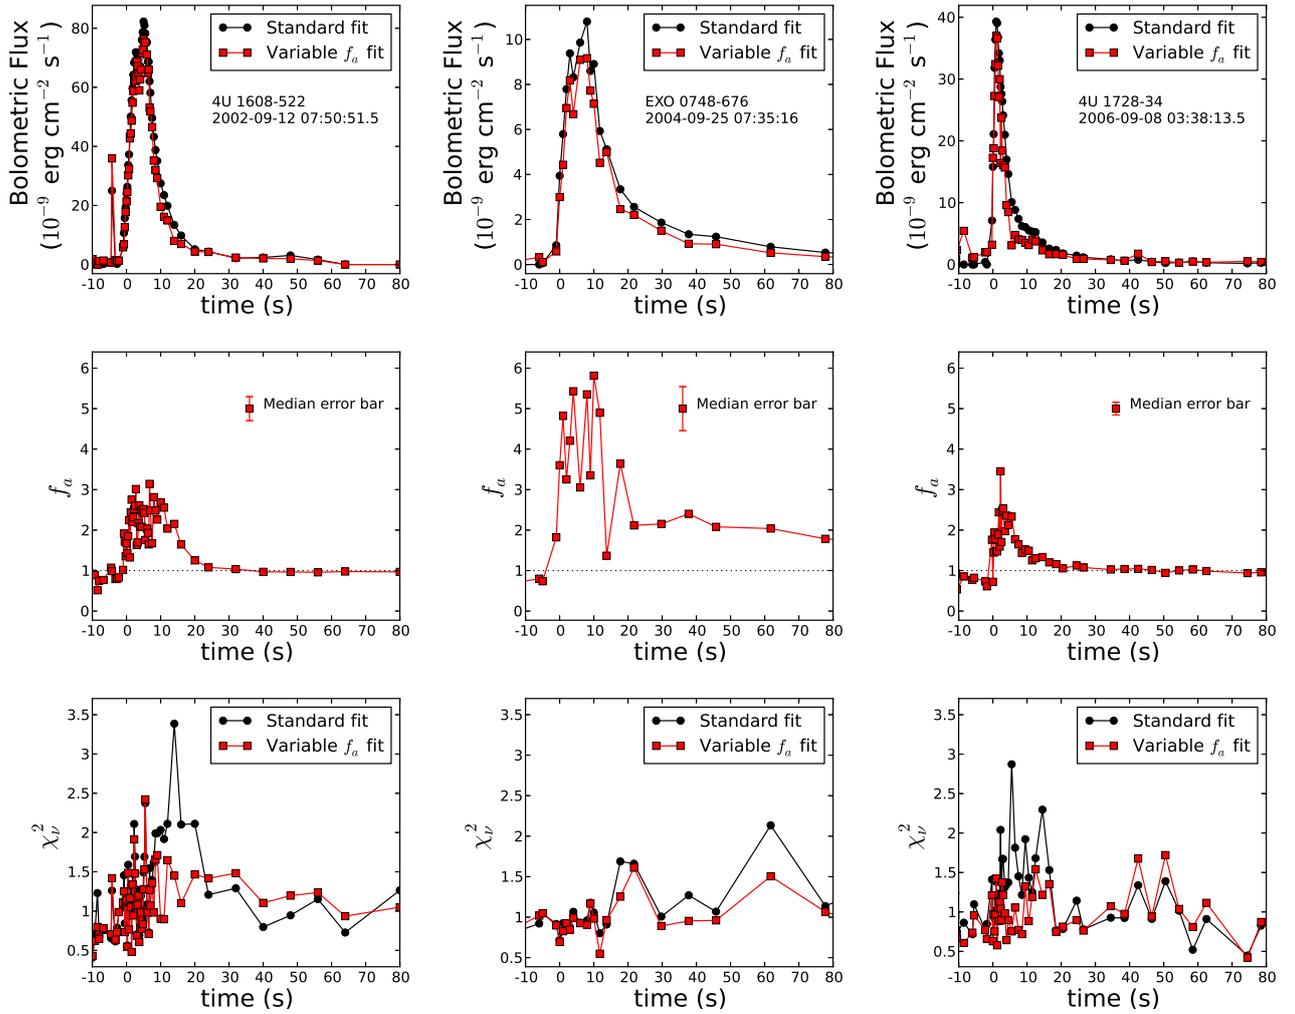


Figure 5.2 Comparison of fitting a variable persistent emission factor, f_a , to the standard approach fits for non-PRE bursts from three different neutron stars. I have selected a typical hydrogen accretor (4U 1608–52), a dipping source (EXO 0748–676), and a He-accretor and likely ultracompact binary (4U 1728–34). The variable fit approach yields consistently lower burst component fluxes (top panels). The contribution to the flux from the variable persistent emission increases to several times the pre-burst level (middle panels) during the rise of the burst. Allowing the persistent emission to vary improves the χ^2_{ν} (lower panels).

closely approximates ratio of likelihoods. As explained in [Gregory \(2005\)](#), there is another contribution to the Bayes factor calculation. Since the variable persistent flux factor f_a represents an additional degree of freedom in the spectral fits, a penalty must be imposed on it. The standard approach can be thought of as a special case of the variable persistent flux approach (with f_a fixed implicitly at 1), so I compared the effective uncertainty for the standard approach (i.e., the Gaussian uncertainties on the pre-burst persistent spectra) with the prior for the variable persistent flux approach, which is flat between $f_a = -100, 100$ as stated in §2.2.2. I then found the median total counts in all the pre-burst spectra, which is approximately 9,000 counts. I adopted 10^4 counts for a characteristic count rate for the pre-burst persistent spectra. This gives a width of the prior in the standard approach of $\delta f_a/f_a = \sqrt{10^4}/10^4 = 0.01$. The overall estimate of the Bayes factor is then

$$B \approx \frac{p(D|M_1, I) \delta f_a}{p(D|M_2, I) \Delta f_a} = \frac{5.77 \times 10^{-3}}{4.51 \times 10^{-9}} \times \frac{0.01}{200} = 64. \quad (5.1.1)$$

Thus I conclude that the variable persistent flux improves on the standard approach with a Bayes factor of 64, even if I take into account the penalty imposed by the extra degree of freedom in the variable persistent flux approach.

5.1.2 Correlation with burst flux

The results of Chapter 3 suggest that, for PRE bursts, there is only a weak correlation between f_a and burst flux. I found that, although most of the high f_a values occurred during radius expansion, there was nonetheless variability in f_a when the burst flux was relatively constant. I have repeated this analysis for non-PRE bursts. I have taken all spectra for each burst between the beginning of the burst and the point at which burst flux declines below 10% of its peak value, and performed Kendall τ rank correlations on burst flux against f_a for each burst. I found positive correlations in 1,000 out of 1,419 non-PRE bursts, of which 919 are at greater than 3σ significance. Performing the same test on the cooling tails of PRE bursts, I found positive correlations in 299 out of 316. Of these, 286 were at greater than 3σ significance. These statistics demonstrate that there is a more direct relationship between burst flux and f_a for non-PRE bursts than for PRE bursts.

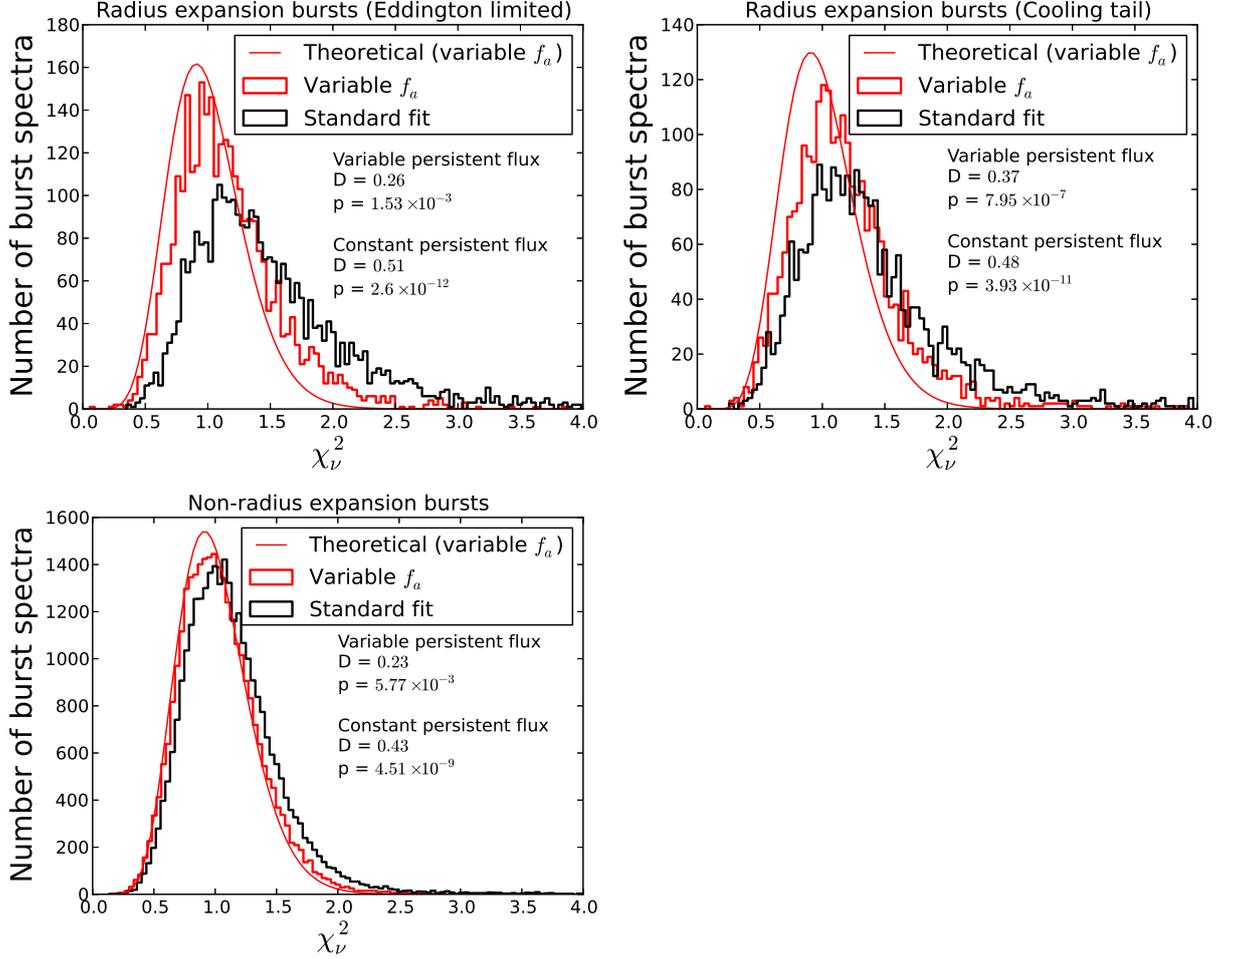


Figure 5.3 Histograms of the χ_ν^2 for the variable persistent normalisation and standard approach spectral fits (red and black histograms respectively), and the theoretical χ_ν^2 distributions for a model that adequately describes the data (red and black curves respectively). The first panel shows spectra from Eddington-limited PRE bursts, the second panel shows cooling tail spectra from PRE bursts, and the third panel shows spectra from non-PRE bursts. It is obvious that the variable persistent normalisation approach gives better fits overall than the standard approach. It can also be seen that a blackbody burst spectrum is far from acceptable in PRE bursts, even in the cooling tail when the atmosphere has returned to the surface of the star. The variable persistent normalisation approach corrects for much of this effect. Non-PRE spectra are better described by a blackbody and constant persistent model, but even these are significantly improved with the variable persistent normalisation approach. Results of Kolmogorov-Smirnov tests are also shown; I list the K-S statistic D and null hypothesis probability p for the variable persistent normalisation and standard fits respectively.

The number of PRE and non-PRE bursts listed in Table B.1 represent the bursts which made it into the final analysis; bursts discarded for source confusion or the inability to fit a persistent emission model are *not* counted. These are still suitable to estimate the Eddington flux of the source. For example, IGR 17473-2721 has three PRE bursts in the complete *RXTE* catalog and therefore has a known F_{Edd} but all three of these bursts are source confused so they are not investigated in this work or listed in Table B.1.

5.1.3 Collation of Eddington fluxes

Table B.1 lists Eddington fluxes for all sources from which at least one PRE burst has been observed by *RXTE*'s PCA or WFC instruments. This data is supplemented by a small number of PRE bursts detected by other X-ray satellites. These events are mostly medium to long duration bursts and are listed in Table B.2 in Appendix B. This sample is heterogeneous, but represents our only knowledge of the Eddington flux of the source SLX 1735–269, where only non-PRE bursts were observed by *RXTE*, and improves the estimates of the Eddington fluxes of the other sources listed.

In Chapter 3 I found that, in PRE bursts, the variations in measured f_a cannot be attributed to counting statistics. I repeated this test for all non-PRE bursts in the catalog. For all spectra obtained during non-PRE bursts I generated 100 simulated spectra using the standard approach spectral parameters, and incorporated counting statistics typical of the detector. I then fit these simulated spectra with the variable persistent normalisation method, obtaining a standard deviation of f_a measurements arising from counting statistics in the absence of any variation of persistent flux. I found f_a measurements greater than unity to 5σ significance in 214 spectra of 62783. In the absence of a significant effect I would expect at most one such case arising by chance. This experiment shows that enhanced persistent emission is present in non-PRE bursts as well as PRE bursts.

Since PRE bursts are by nature brighter than non-PRE bursts, it is possible that the difference in χ^2_ν between the two classes is simply due to superior signal-to-noise in the brighter bursts. To investigate this I took a subset of spectra from both PRE and non-PRE bursts, with bolometric fluxes between 4.0×10^{-8} and 5.0×10^{-8} erg s⁻¹ cm⁻² and blackbody temperatures between 2.0 and 2.5 keV, these quantities being measured with

the standard approach. There were 331 spectra from PRE bursts and 655 spectra from non-PRE bursts in this segment of the parameter space. These had mean χ_ν^2 of 1.17 and 1.35 for non-PRE and PRE bursts respectively, and a K-S test indicated a probability of 1.79×10^{-6} that these χ_ν^2 are drawn from the same distribution. This shows that the poorer spectral fits generally observed in PRE bursts are not due to better signal-to-noise in the bright bursts.

5.2 Is the persistent emission spectral shape constant?

The approach of treating the persistent emission as a contribution of fixed shape depends upon the assumption that its shape does not change, at least not on timescales of the order of the burst duration, or to a degree sufficient to affect the overall spectral fits. I did not investigate this matter directly in Chapter 3, but my finding that f_a generally returns to its pre-burst level, with good χ_ν^2 , suggests that this assumption is not unreasonable. Similarly, [Keek et al. \(2014\)](#) and [Peille et al. \(2014\)](#) also report that the persistent emission returns to its original level soon after the burst. [Thompson et al. \(2005\)](#) and [Bagnoli et al. \(2013\)](#) have found that the persistent emissions of 4U 1826–24 and the Rapid Burster respectively are remarkably stable over kilosecond timescales between bursts. [Linares et al. \(2014\)](#) found variability over ~ 10 minute periods in the *intensity* of the persistent emission of IGR J18245–2452, with no change in the *shape* of the persistent emission. It is therefore likely that no change in the shape or intensity of the persistent emission is to be expected in the absence of burst luminosity.

5.2.1 Short-term variability

To determine whether the persistent emission shape and intensity are significantly variable on shorter timescales than minutes or hours, in the absence of a burst, I took every 0.25s pre-burst spectrum prior to every burst and fit it with just f_a times the persistent model for that burst, i.e., no thermonuclear burst emission. The distributions of f_a and χ_ν^2 are shown in Figure 5.4. This data typically covers a few tens of seconds prior to every burst. It is plain to see that f_a is distributed tightly around unity, indicating that the persistent emission generally does not vary greatly in intensity on timescales of 1s or less.

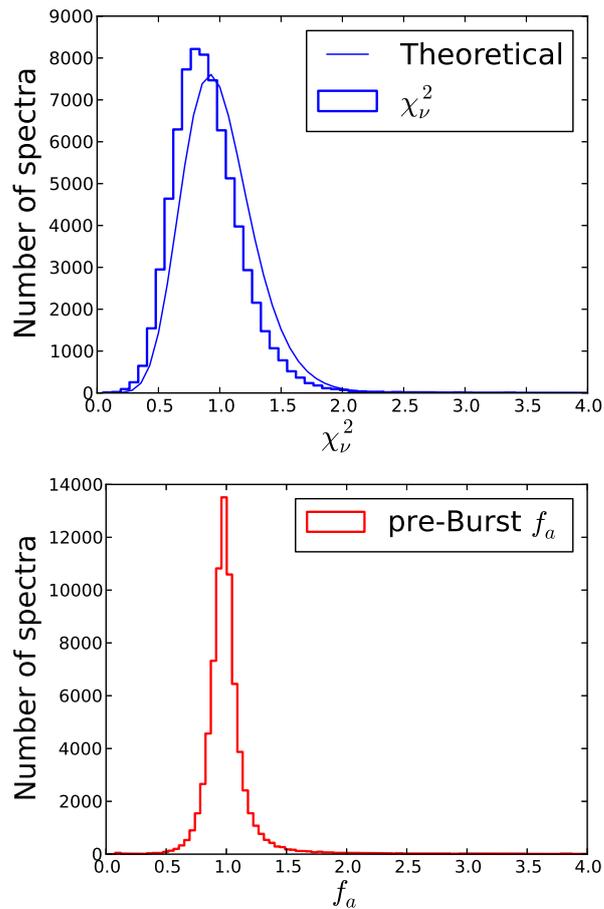


Figure 5.4 Measured distribution of f_a (red histogram), and measured and theoretical χ_ν^2 (thick and thin blue curves) for pre-burst emission in the absence of any nuclear burning. The f_a values are very strongly peaked around 1, and the measured χ_ν^2 are low, indicating that the pre-burst emission does not change intensity or spectral shape to any detectable degree.

The χ_ν^2 distribution also implies that the shape of the persistent emission is not changing appreciably, although of course the short exposure times and associated low photon counts make detecting such changes difficult.

5.2.2 Medium-term variability

To test for more subtle spectral shape changes, I increased the exposure time to obtain better signal-to-noise. For the three bursts studied in Figure 5.2, and the fifteen bursts investigated by (Peille et al. 2014; see their Table 1), I inspected the *RXTE* Standard-2 data covering the entire observation. These spectra are binned into 16s intervals. I fit each of these spectra with the persistent model for that burst, with all parameters frozen, and f_a normalisation allowed to vary. I considered only spectra recorded within an hour before

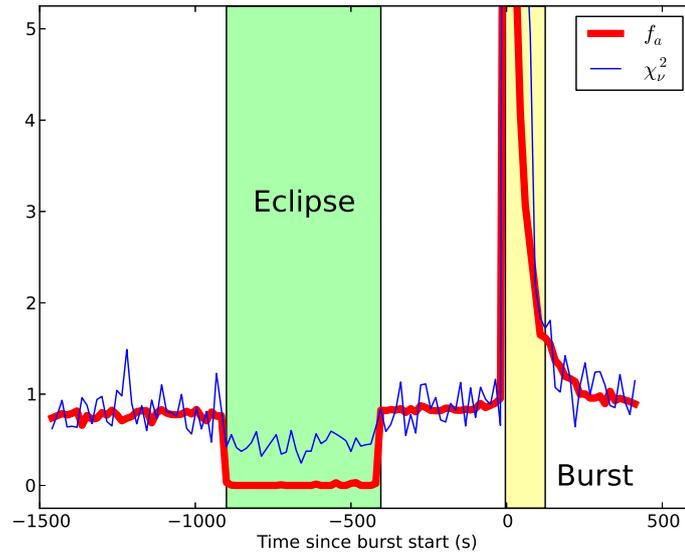


Figure 5.5 Fits to the persistent emission just prior to and just after a burst from EXO 0748–676, using only variable persistent flux and no burst blackbody component. The shaded areas (an eclipse by the companion star, and the burst itself) are discarded. It is evident that f_a remains approximately constant over the observation. The χ^2_ν also does not change, suggesting that the spectral shape of the persistent emission does not vary significantly.

or after the burst but excluding those *during* the burst, to avoid contamination by the burst emission. That is, I excluded spectra recorded between the beginning of the burst (burst flux has reached 1/4 of the peak burst flux) and the end of the PCA data of the G08 catalog. This data typically covers a few minutes to an hour before and after the burst. These features are shown in more detail in Figure 5.5. I also excluded an eclipse of EXO 0748–676 beginning at 2004-09-25 07:20:29, 15 minutes before the burst, when no flux was observed from the neutron star (see also Homan et al. 2003). I analysed the distribution of f_a and χ^2_ν . These are shown in Figure 5.6. These results suggest that the persistent emission does not change greatly in either shape or intensity on time scales up to an hour before and after a type I burst.

5.2.3 Detectability of persistent emission spectral shape change

Finally, I investigated the extent to which I can detect changes in the accretion spectral shape if they really are present, assuming that a change in the intensity of the persistent emission reflects a change in the accretion rate. I selected a non-PRE burst from 4U 1636–536 that

uses the `disko` model (Stella & Rosner 1984), recorded 2005 Apr 4, 10:48:43. For this burst PCUs 0 and 2 on *RXTE* were active. This model contains $\dot{M}/\dot{M}_{\text{Edd}}$ explicitly as a variable parameter and incorporates physically-motivated changes in the spectral shape; for this burst the best fit accretion rate was $0.205\dot{M}/\dot{M}_{\text{Edd}}$ and the peak f_a was 3.3. The normalisation of this persistent model was $2 \cos i/d^2 = 1.648$, where i is the inclination and d is the distance in units of 10 kpc. At an assumed distance of 6 kpc (Galloway et al. 2006), this gives an inclination of 72° , in agreement with Pandel et al. (2008) and Casares et al. (2006), who also infer a high inclination (greater than 64° , and $36^\circ - 74^\circ$ respectively) for this system. The normalisation was kept frozen in the subsequent analysis. I generated 2,000 simulated spectra, folded through a PCA response appropriate to the observed burst, of the `wabs*disko` model with \dot{M} taking random values between $0.067\dot{M}_{\text{Edd}}$ and $0.667\dot{M}_{\text{Edd}}$. I then fitted those simulated spectra with f_a times the original persistent model. The measured f_a and χ_ν^2 are given in Figure 5.7, left two panels. In the absence of burst emission I can clearly detect a change in spectral shape via the increase in χ_ν^2 with increasing \dot{M} as well as an increase in the intensity of the persistent emission. I then repeated the process, this time including a blackbody with temperature 2.0 keV and normalisation representing a sphere with radius 10 km at a distance of 6 kpc (i.e., the surface of 4U 1636–536). This I fit with the usual variable persistent flux approach (right two panels). It is evident that, although I can still detect enhanced persistent emission through an increased f_a , it is no longer possible to detect a change in its spectral shape. This experiment gives confidence that the approach of simply varying the normalisation of the persistent emission will introduce no detectable systematic effects.

5.2.4 Variability outside the 2.5-20 keV energy range

This analysis is not sensitive to changes in the persistent spectrum outside the 2.5-20 keV energy range I consider in this thesis. In't Zand et al. (2013) combined data from the *Chandra* X-ray satellite, which is sensitive at energies below 2keV, with *RXTE* data for a radius expansion burst from SAX J1808.43658. They obtained similar f_a values as I did studying the same burst with only *RXTE* data, suggesting that changes in the persistent spectrum below 2.5keV are negligible. At high energies, there is evidence that the hard X-ray (above about 30 keV) flux decreases during a burst, in apparent conflict with the

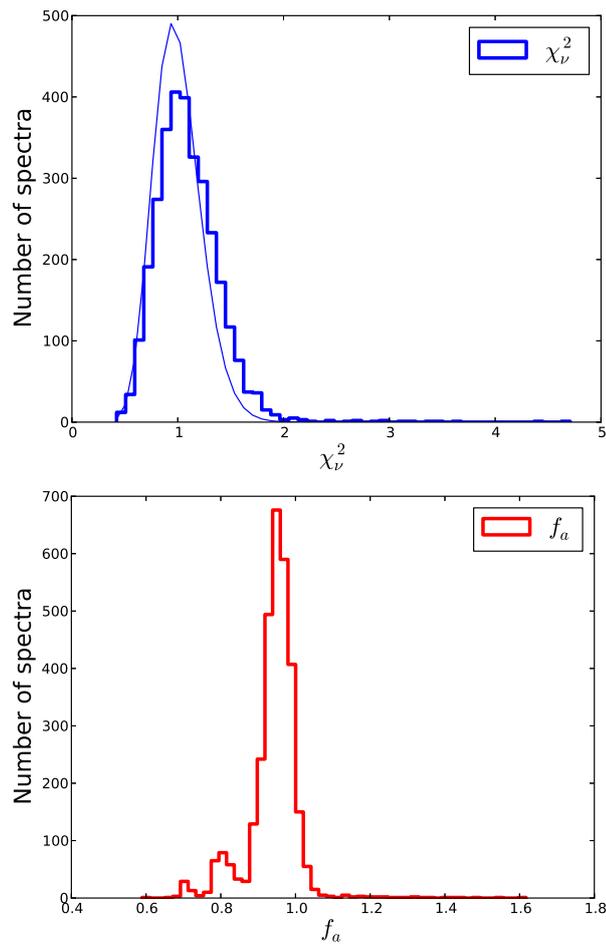


Figure 5.6 Measured distribution of f_a (red histogram), and measured and theoretical χ_ν^2 (thick and thin blue curves) for Standard-2 data covering the observations of eighteen selected bursts. The f_a values are tightly clustered around unity, and the χ_ν^2 are somewhat lower than the theoretical distribution assuming a well-fitting model. This suggests that the persistent emission does not change shape or intensity to any detectable degree on timescales of an hour before and after a type I burst.

results of this work (Chen et al. 2013, Ji et al. 2013, Ji, Zhang, Chen, Zhang, Torres, Kretschmar & Li 2014, Ji, Zhang, Chen, Zhang, Kretschmar, Wang & Li 2014). However, this phenomenon is attributed to the rapid cooling of the accretion disc corona, which contributes the majority of the very hard X-rays. An increase in 2.5-20 keV flux attributed to accretion rate change does not conflict with a simultaneous quenching of high energy photons through cooling of the corona.

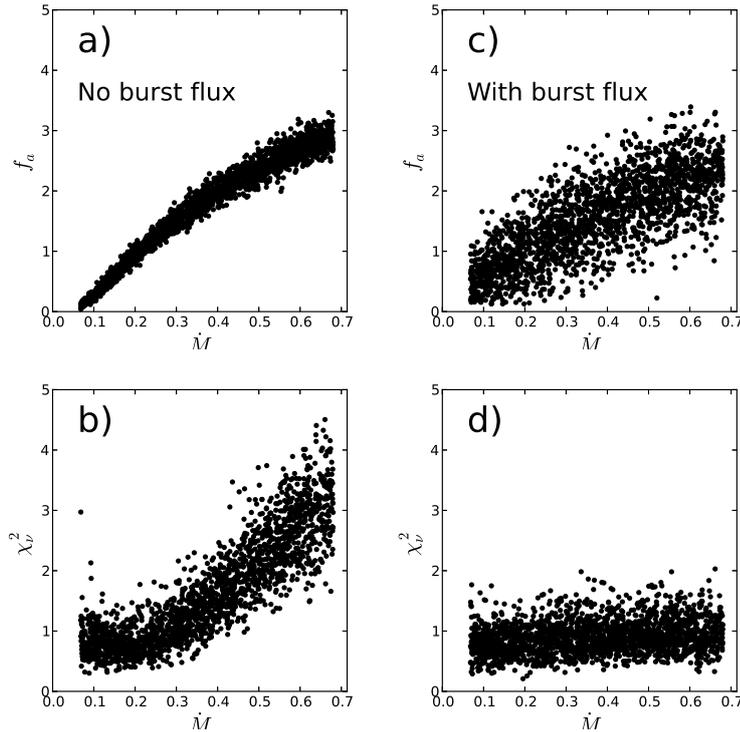


Figure 5.7 Results of fits to simulated *disko* spectra with \dot{M} allowed to vary, both with (panels a. and b.) and without (panels c. and d.) a burst component. In both cases an increase in f_a is detected. When there is no burst component, and only a change in the intensity of the persistent emission is being tested for, the increasing χ^2_ν with increasing \dot{M} indicates that changes in the spectral shape are detectable. In the presence of burst emission, it is no longer possible to discern the spectral shape changes in the persistent emission. I can therefore simply vary the normalisation of the persistent emission without introducing systematic effects.

5.3 Eddington fluxes from non-PRE bursts

For neutron stars for which no radius expansion bursts have been observed, the only lower bound on their Eddington fluxes are the peak fluxes of their brightest bursts. A better knowledge of their Eddington fluxes would be of observational utility, as many of these sources show interesting properties: for example, the prolific burster GS 1826-24 exhibits very regular bursts that are in close agreement with numerical models of X-ray bursts (e.g., Heger, Cumming, Galloway & Woosley 2007, Galloway & Lampe 2012). The 11 Hz pulsar IGR 17480–2446 is the most slowly rotating source that shows burst oscillations (Cavecchi et al. 2011), making that source an important test of burst oscillation models. Neither source has a single reported PRE burst. The Rapid Burster (MXB 1730-335) is the major source of knowledge about type II bursts, and has one reported PRE burst (Sala

et al. 2012). This event, however, was detected by the *Swift* X-ray instrument and the measurement was not sensitive enough to constrain the Eddington flux to better than about 50% (see their Fig. 2).

An interesting observation is that all the f_a results of Chapter 3 fall below the line $f_a\gamma = 1$, where γ is the pre-burst accretion flux as a fraction of the Eddington flux, consistent with predictions of Burger & Katz (1983) and Miller & Lamb (1996). This observation raises the possibility of obtaining a lower bound on the Eddington flux for a given source, that is, $F_{\text{Edd}} > f_a F_{\text{Pers}}$, where F_{Pers} is the bolometric flux of the pre-burst persistent emission.

In Figure 5.8 I show peak f_a against γ for every type I burst detected by RXTE. For $0.01 < \gamma < 0.4$ the non-PRE bursts appear to be bounded above by $f_a\gamma \lesssim k$, where k appears to be somewhat less than 0.5. This would give a lower limit on the Eddington flux roughly twice as constraining as the most conservative estimate $f_a F_{\text{Pers}} < F_{\text{Edd}}$. Determining γ requires knowing F_{Edd} ; the method for determining these for PRE bursters is given in §2.3.

5.3.1 Estimating γf_a upper limit

I now describe a method by which I estimate the upper bound of $f_a\gamma$ that is reproducible and more objective than estimating by eye. I treat the measured $x_i = (f_a\gamma)_i$ values as normal distributions centered around their measured values, with standard deviations σ_i equal to the uncertainty in the f_a , times γ . This normal distribution is given by

$$G(x, \sigma_i, x_i) = \frac{1}{\sqrt{2\pi}\sigma_i} \exp\left[-\frac{(x - x_i)^2}{2\sigma_i^2}\right]. \quad (5.3.1)$$

The fraction of this distribution lying above x_{max} is

$$\begin{aligned} A(\sigma_i, x_i) &= \int_{x_{max}}^{\infty} G(x, \sigma_i, x_i) dx \\ &= \frac{1}{2} \left[1 + \operatorname{erf}\left(\frac{x_i - x_{max}}{\sqrt{2}\sigma_i}\right) \right], \end{aligned} \quad (5.3.2)$$

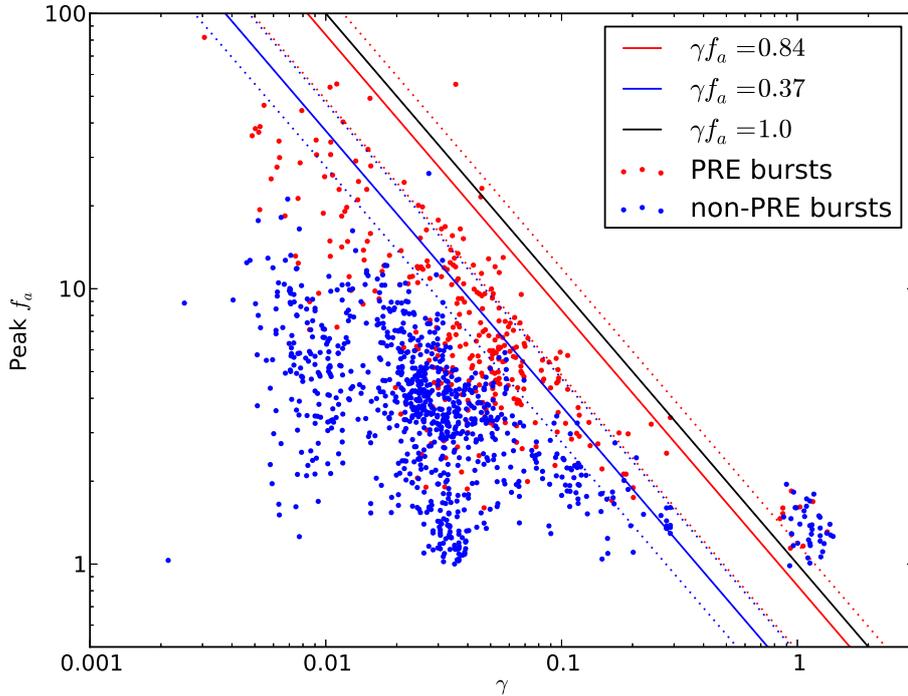


Figure 5.8 Peak f_a against γ for every burst in the sample from sources for which Eddington fluxes are listed in B.1. PRE bursts are shown in red, and non-PRE bursts are shown in blue. There is a clear division between the two classes of bursts. PRE bursts are bounded above by $f_a\gamma \lesssim 0.84$, while the non-PRE bursts are bounded above by $f_a\gamma \lesssim 0.37$ (as calculated by the method explained in §5.3). The dotted lines indicate 1σ intervals estimated using the bootstrap method. I also show the hard upper bound $f_a\gamma < 1$ (e.g. Burger & Katz 1983). The clump of points at $\gamma \approx 1$ represents bursts from the rapid accretors Cyg X-2 and GX 17+2; these are atypical bursters and I do not consider their bursts in deriving $f_a\gamma$ relations.

and so the total over all data points is

$$A_T = \frac{1}{2} \sum_{i=1}^n 1 + \operatorname{erf} \left[\frac{x_i - x_{max}}{\sqrt{2}\sigma_i} \right], \quad (5.3.3)$$

where n is the number of data points.

The question is how to find x_{max} . First I make a guess ζ as to the value of x_{max} and calculate a TA from this and the data points x_i . Then I take a set of n data points \hat{x}_i with random values between 0 and ζ but the measured uncertainties σ_i ; this step assumes the observed data points are evenly distributed with uncertainties independent of their values, but the procedure can be modified to suit other assumptions. From this random set \hat{x}_i I calculate an expected total area ($TA(rand)$) that can be compared to the real one

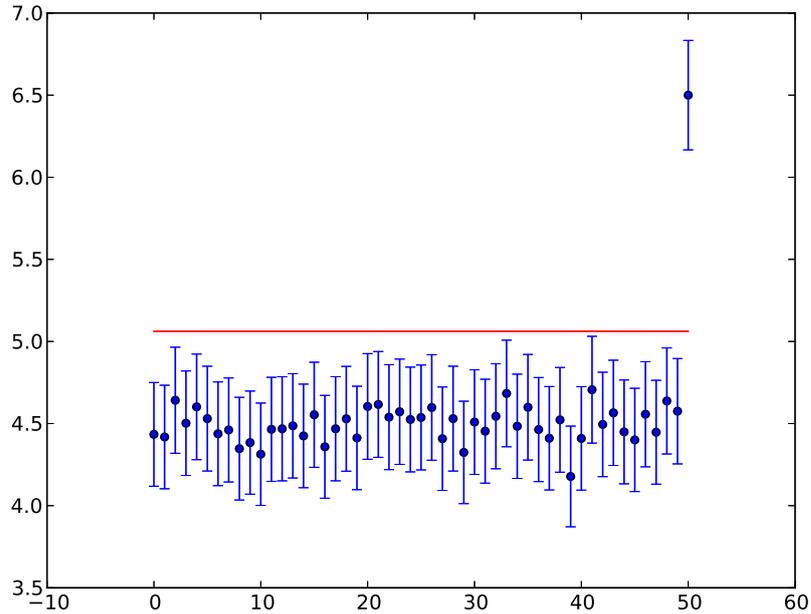


Figure 5.9 A set of 50 randomly placed points and one extreme outlier (blue points) and the estimated upper bound according to the method developed in this section (red line). It is clear that this method gives a good estimate for the upper bound of the data points without assigning too much weight to the outlier. The units on both axes are completely arbitrary.

($TA(\text{observed})$). Making many realisations of the random set, I get an expected $TA(\zeta)$ and a standard deviation SD on that. To avoid the analysis assigning too much weight to outliers, I take many bootstrap subsamples of the original data set and take the overall x_{max} to be the average of those from the bootstrap samples; this procedure also gives us uncertainties on x_{max} . In Figure 5.9 I demonstrate this method for a set of randomly placed points containing one outlier.

Application to sources with known Eddington flux

Applying this procedure, I get $x_{max} = 0.84$ for the PRE bursts and $x_{max} = 0.37$ for the non-PRE bursts. I note that the $f_a\gamma$ relation is a phenomenological one and does not depend on any physical interpretation, though it is consistent with the requirement that the accretion luminosity not exceed Eddington. To test whether the new method gives a more stringent lower bound than the peak burst flux, I compared the two methods for every source for which ten or more non-PRE bursts have been observed, but excluding the rapid accretors Cyg X-2 and GX+17-2. I also exclude SLX 1744-300, as that source is

difficult to distinguish from the nearby ($\Delta\theta \approx 3'$) SLX 1744–299 (Skinner et al. 1990). I first apply the method to sources for which the Eddington flux is known, to make sure the new method never overestimates the Eddington flux. In 2 of 16 sources, the new method gave a greater lower bound than the brightest non-PRE burst, and in no case did the new method overestimate the Eddington flux beyond observational uncertainties. These results are shown in Figure 5.10.

Application to sources with unknown Eddington flux

Applying the new method (with $x_i = (f_a F_{\text{Pers}})_i$ since γ is unknown) to sources with unknown Eddington flux, I get a better bound than the non-PRE bursts for the sources Cir X–1 and IGR 17480–2446, with F_{Edd} lower limits 18×10^{-9} and 42×10^{-9} erg s $^{-1}$ cm $^{-2}$ respectively. IGR 17480–2446 is located in the globular cluster Terzan 5, with a known distance of about 5.5 kpc (e.g., Papitto et al. 2011), and a hydrogen accreting neutron star with $1.4M_{\odot}$ would have an Eddington flux of 57×10^{-9} erg s $^{-1}$ cm $^{-2}$. Cir X–1 is believed to be 7.8 to 10.5 kpc distant (Jonker & Nelemans 2004). This result suggests it is probably at the nearer end of that range, since the more distant value would give an Eddington flux of 16×10^{-9} erg s $^{-1}$ cm $^{-2}$ for a canonical neutron star, slightly higher than the estimate given here. I also get a marginally more constraining Eddington flux for 1M 0836–425 of 23×10^{-9} erg s $^{-1}$ cm $^{-2}$ than the brightest burst detected from that source.

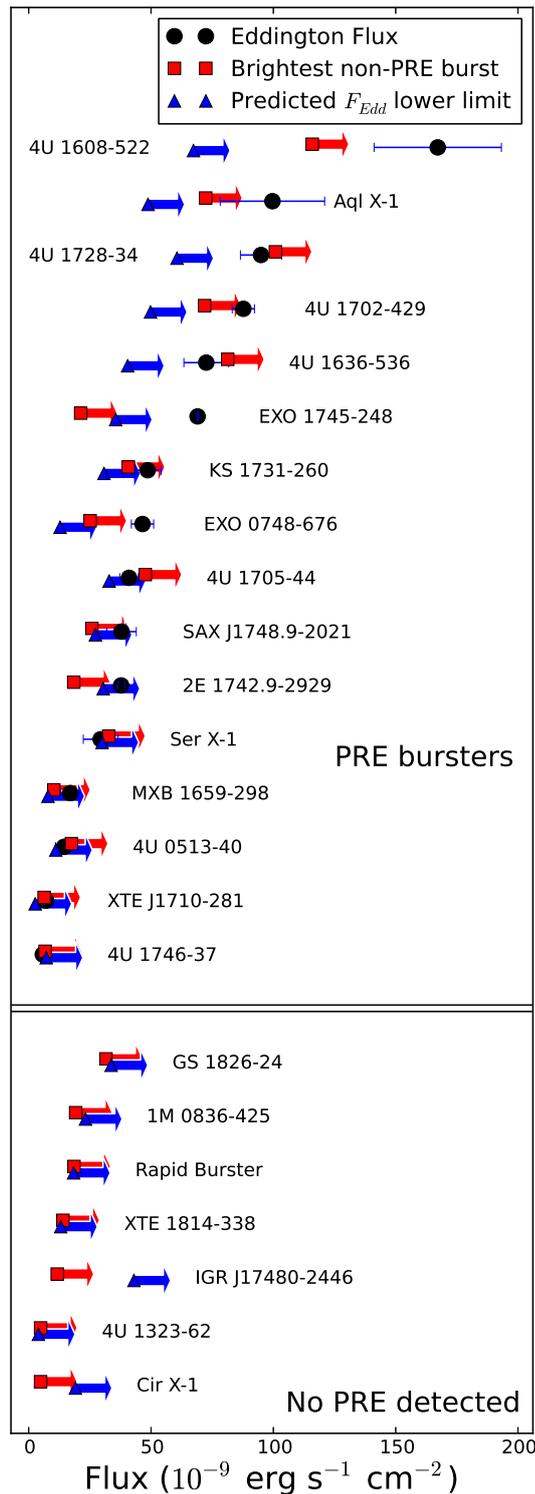


Figure 5.10 Lower limits of F_{Edd} for various neutron stars from their non-PRE bursts using the method described in §5.3. I have plotted the brightest non-PRE burst (squares) and the new prediction (triangles). Right-pointing arrows indicate that these are lower bounds on the Eddington flux. For sources with known F_{Edd} these are shown (circles) to demonstrate that the new method does not overestimate the Eddington flux. For sources where F_{Edd} has not been directly measured, I have obtained a significantly better constraint for the F_{Edd} of Cir X-1 and IGR 17480–2446, and a marginally better constraint for 1M 0836–425, than those given by their brightest non-PRE bursts.

Chapter 6

SPH simulation setup and results

On two occasions, I have been asked, “Pray, Mr. Babbage, if you put into the machine wrong figures, will the right answers come out?” I am not able rightly to apprehend the kind of confusion of ideas that could provoke such a question.

Charles Babbage

The results of previous Chapters suggest that the accretion rate can be substantially increased during a burst, if the measured increase in persistent emission reflects an increase in mass transfer rate. In this Chapter I use three-dimensional SPH simulations to test the hypothesis that radiation drag induced by luminosity at the stellar surface can cause such an effect.

I have modelled the simulated burst light curve after a typical observed burst (see §6.3), and select a disc configuration (see §6.1.2) that produces a similar pre-burst accretion rate to that observed for this burst. Radiation transport through the disc is modelled with a range of descriptions chosen to be simple to calculate and which approximate physically plausible behaviour.

6.1 Simulation Setup

6.1.1 Neutron star

I take the central neutron star to have “canonical” values of mass $M_* = 1.4M_\odot$ and radius $R_*=10$ km. In my PHANTOM simulation the star is represented by a gravitational potential at the origin, and an accretion radius of 10 km also centered on the origin. These are fixed at the origin. SPH particles crossing the accretion boundary are deemed to be accreted onto the star, liberating their mass as energy as described in §6.4.3.

6.1.2 Disc

I only simulate the inner part of the accretion disc. Since the disc extends to the neutron star surface, where the Keplerian orbital speed is very fast, 45% of the speed of light, the computational time step is very small ($dt < 1.0 \times 10^{-6}$ s). This orbital period at the neutron star surface equates to a frequency of 2,170 Hz, which is fast compared to neutron star spin periods of less than 700 Hz (see A.1). It would be prohibitive to simulate the entire extent of the disc and so I let the outer edge of the simulation be at 8,000 km. At this radius the Keplerian orbital period is about ten seconds, the length of the burst.

The disc is given an initial mass of M_{disc} and surface density profile

$$\Sigma(R) = \sqrt{R_*} M_{\text{disc}} R^{-3/2} \left(1 - \sqrt{\frac{R_*}{R}} \right), \quad (6.1.1)$$

where R is the cylindrical radius from the center of the neutron star. The final term is there to ensure the disc has zero surface density at the surface of the star.

According to Frank et al. (2002) the scale height of an irradiated disc should satisfy

$$H(R) \propto R^{9/7}. \quad (6.1.2)$$

This also ensures that the disc is convex. Assuming that $c_s \propto R^{-q}$ for the sound speed and $\Omega \propto R^{-3/2}$, the Keplerian orbital velocity, then solving for q gives $q = 3/14$. For the

constant of proportionality I take $H(R_*) = 0.02R_*$. Furthermore,

$$H = c_s \frac{R^{3/2}}{(GM_*)^{1/2}}. \quad (6.1.3)$$

But since $c_s = c_{s,0}R^{-q}$ I have

$$H = c_{s,0} \frac{R^{3/2-q}}{(GM_*)^{1/2}}, \quad (6.1.4)$$

and $c_{s,0}$ is determined by the H/R set at the inner boundary.

The mass of the disc is as yet unspecified. It is, however, known to be negligible compared to the mass of the neutron star. Self-gravity can be neglected, meaning that SPH particles will behave the same way under the influence of viscosity and gravity regardless of the total mass of the disc. The only role disc mass will play is in the calculation of opacity, the mass accretion rate, and the luminosity generated by accretion onto the star. Thus I can find M_{disc} by initially setting the mass to some arbitrary value and running a simulation with no irradiation. Given some target mass accretion rate (I choose 2% of \dot{M}_{Edd}) I can then adjust the disc mass to achieve this.

I caution that the idea of a single well-defined Eddington luminosity is illusory due to the dependence of opacity on temperature (see Equation 1.3.4; Lewin et al. 1993, eq. 4.13b). The Eddington luminosity for material at a temperature of 3 keV is 10% lower than that of cold material. The Eddington luminosity can therefore be expected to increase in the disc with distance from the neutron star. The disc itself has a temperature of 0.5 to 1.0 keV (Paizis et al. 2006). If, as suggested by Walker (1988), interactions between disc and radiation drive electrical currents, there might be additional localised opacity changes associated with enhancement or depletion of electrons. In this work I do not attempt to account for such effects and take a constant, frequency independent opacity $\kappa = 0.3156 \text{ cm}^2 \text{ g}^{-1}$, the opacity of material with hydrogen fraction $X = 0.7$ and temperature $T = 2.0 \text{ keV}$. This ensures a single, well-defined Eddington luminosity everywhere.

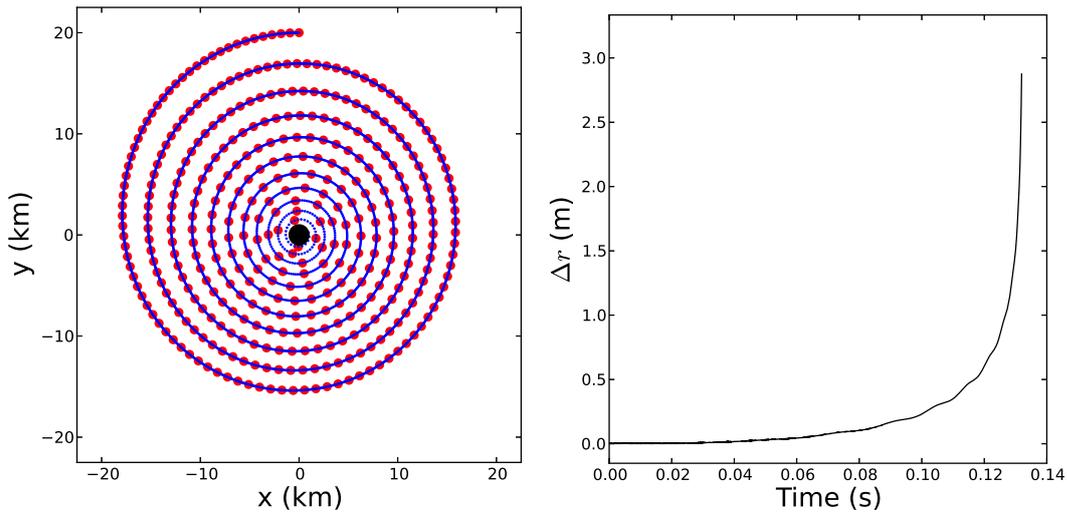


Figure 6.1 The left panel shows the trajectory of a test particle calculated using PHANTOM (blue curve) and the test code (red curve). The particle was placed 200 km from a $1.4M_{\odot}$ neutron star (black circle of radius 10 km), given initial v_{ϕ} of 0.9 times the Keplerian orbital speed, and experiences radiation forces of 10% Eddington. For clarity, only some of the points are plotted. The codes are indistinguishable on this scale. In the right panel is shown the distance (Δr) between the PHANTOM particle and the test code particle. At the moment of impact with the neutron star, the particles are only 3 m apart. These figures show that `extern_prdrag` is handling the equations of motion correctly.

6.2 Numerical tests

6.2.1 Correct integration of equations of motion

To test the correctness of the code, I wrote a new test code that does nothing but integrate the equation of motion 3.4.1, using the Euler method, for a single particle, a fixed β , and some specified initial position and velocity. I ran this code for a $1.4M_{\odot}$ neutron star with $\beta = 0.1$ and a particle orbiting at $R = 200$ km with initial v_{ϕ} of 0.9 times the Keplerian orbital speed. I then ran a PHANTOM setup for a spherical assembly of 89 nonabsorbing and nonviscous particles centered at the same location and with the same initial velocity. Although it is possible to perform an SPH simulation with just one test particle, I would have had to turn off density calculations in PHANTOM, and it was easier to use a collection of nearly coincident particles instead. The radius of the spherical assembly was very small, 1 cm in radius, to avoid the assembly stretching and flattening along its trajectory. As shown in Figure 6.1, there is no significant difference between the codes. I therefore expect that the behaviour of the SPH disc particles under the influence of some β will be correct.

6.2.2 Resolution of scale height

As the neutron star disc is geometrically thin, it is necessary to have enough particles to resolve the vertical height of the disc. I set up several unilluminated simulations with inner and outer radii of 10 km and 100 km respectively, with p - and q indices of 3/2 and 3/4 respectively. This choice of indices ensures that the ratio of scale height to particle density is approximately constant throughout the disc. In simulations of the full disc, the q -index is 3/14, i.e., the ratio of scale height to particle density increases with distance from the neutron star. Thus the scale height in the full simulation is actually better resolved than these test simulations.

I also ran a 1D code that solves the diffusion equation 1.6.12 (e.g., Pringle 1981), with the same physical setup. I truncate the disc at 100 km for the PHANTOM simulations, but impose a rapid dropoff at this radius for the 1D code, as it cannot handle sharp discontinuities or grid points with $\Sigma = 0$. H/R at the inner edge of the disc is 0.02. I ran several of these simulations, with different numbers of particles: 20,000, 50,000, 100,000, 200,000, 500,000, 1,000,000 and 2,000,000. I ran the simulations for 1,000 code units of time (0.03 s of physical time) and compared them with each other. Results are shown in Fig 6.2.

I found that the simulations converged with 100,000 particles. Only minor improvements are evident after that, and these cause significant increases in simulation run time. It follows that, to get acceptable results, I require 100,000 particles between 10 km and 100 km. This consideration makes it possible to calculate the number of particles required to model the entire disc to 8,000 km: the total disc mass out to a radius r_{out} is

$$\begin{aligned} M(r_{\text{out}}) &= \int_{r_{\text{in}}}^{r_{\text{out}}} 2\pi r \Sigma_0 r^{-3/2} \left(1 - \sqrt{r_{\text{in}}/r}\right) dr \\ &= 2\pi \Sigma_0 \left[2\sqrt{r_{\text{out}}} - \log_e r_{\text{out}} - 2\sqrt{r_{\text{in}}} + \log_e r_{\text{in}}\right], \end{aligned} \quad (6.2.1)$$

from which it follows that the number of particles needed is about 2.5×10^6 .

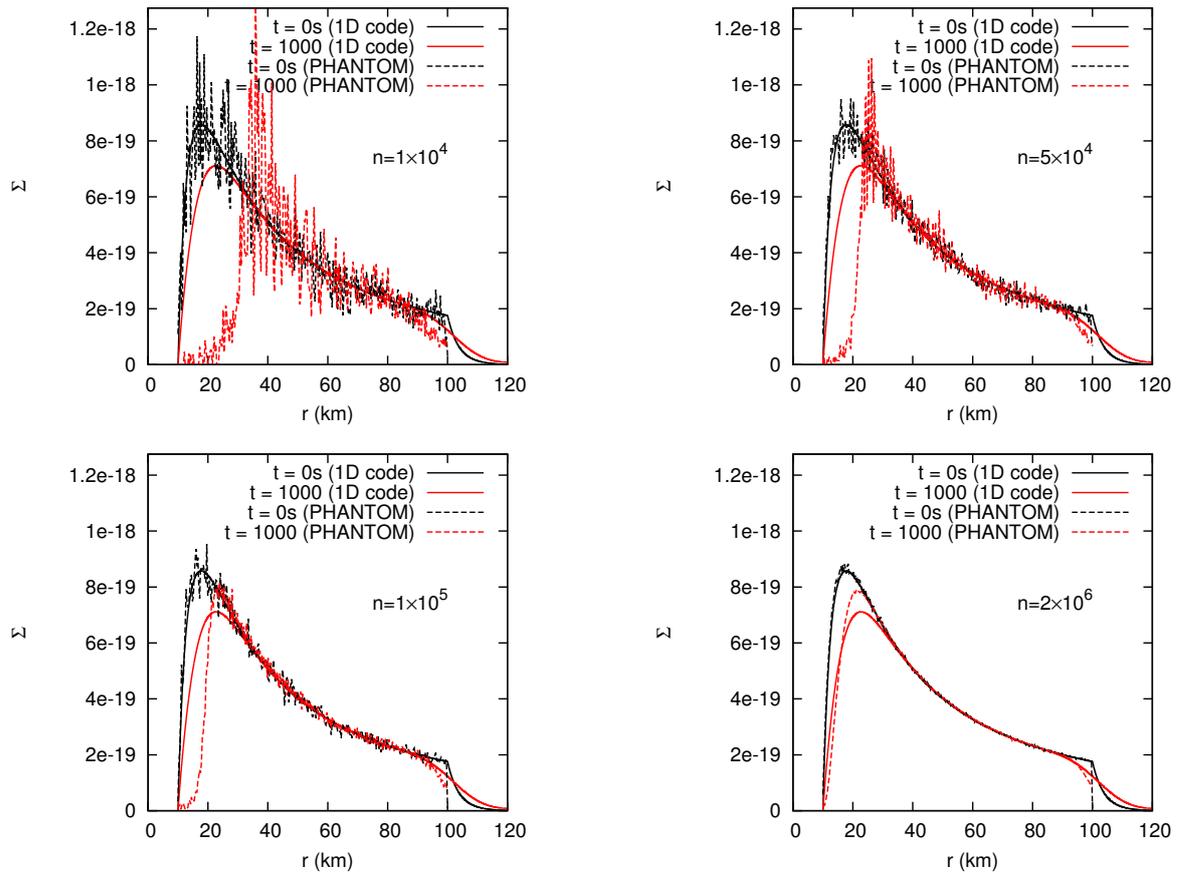


Figure 6.2 Comparison between 1D and 3D codes. I see that the initial setups agree very well regardless of the number of particles in the 3D simulations. After $t=1,000$, the codes are in good agreement for particle counts of 100,000 or higher. The units of surface density are arbitrary. I truncate the disc at 100km for the PHANTOM simulations but impose a smoother dropoff for the 1D code; these features are plainly visible.

6.2.3 Convergence of accretion rate

I then ran simulations of the full disc (10 km to 8,000 km) with different particle counts to ensure that accretion rates also converge. These simulations were run for 60,000 time units (2.0 s). The results are shown in Figure 6.3. It can clearly be seen that the accretion rates have converged to around 2.5% of the Eddington rate, well before the 2.5 million particles I need to resolve the scale height of the disc.

The accretion rate at the beginning is very high and rapidly drops off. This initial transient is mostly due to the particles, initially placed randomly, moving into a more ordered configuration. Another contribution is the fact that the center of gravity of a random distribution of particles will not be exactly at the origin. Both effects tend to cause particles

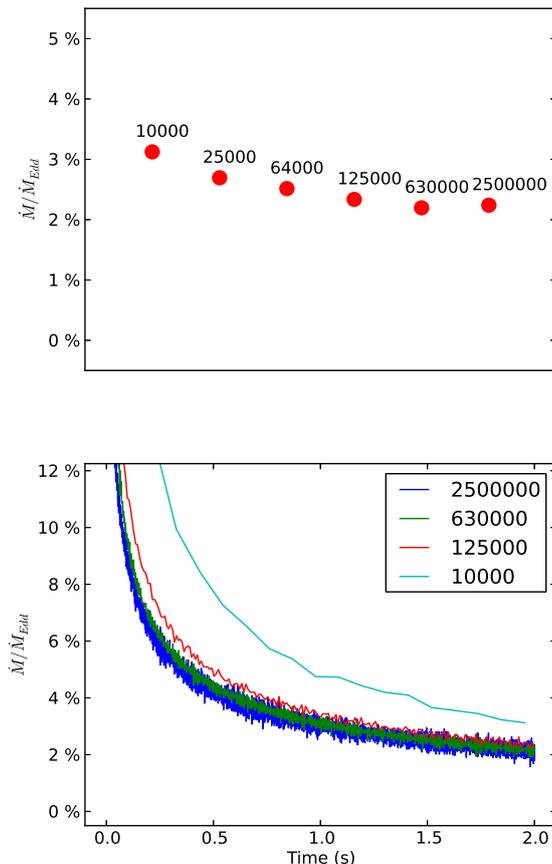


Figure 6.3 Accretion rates for different particle counts. The top panel shows the accretion rate after 2 s for all simulations. The bottom panel shows the instantaneous accretion rates for four of the simulations. They have clearly converged for 630,000 particles, less than the 2.5 million necessary to resolve the scale height of the disc. The bottom panel clearly also shows the need to settle the initial transient.

initially placed near the star to be accreted. Once the initial transient has settled the disc evolves in a stable way, with the accretion rate declining steadily as the disc depletes.

6.3 Burst profile

I have chosen a typical type I burst from the prolific burster 4U 1636–536 as a model burst for the simulations. The burst was observed on 14 Jun 2001, 16:36 and is not a radius expansion burst. I have approximated the light curve by the equation

$$\frac{L(t)}{L_{\text{Edd}}} = \begin{cases} 0, & t_b \leq 0 \\ 9.4t_b^{1.6} \exp[-2.9t_b^{0.4}], & t_b > 0 \end{cases}, \quad (6.3.1)$$

where $t_b = t - 3$, i.e., three seconds after the beginning of the simulation. The first two seconds of the simulation have zero luminosity and are dedicated to the removal of the initial transient (see §6.2.3. The next second of time has zero burst luminosity, but I gradually introduce luminosity caused by the accretion of material as explained in §6.4.3. The observed burst and approximated burst profile are shown in Figure 6.4. According to the catalogue of G08, the star was accreting at approximately 3.1% of the Eddington accretion rate prior to the burst. As discussed in §6.2 the simulated disc accretes at a slightly lower rate (2.5% \dot{M}_{Edd}) after the removal of the initial transient. Because I do not introduce new SPH particles to replace the ones that are accreted the disc depletes over time and the accretion rate slowly declines. It is therefore difficult to measure a pre-burst accretion rate objectively and so the difference between 2.5% and 3.1% \dot{M}_{Edd} is unimportant.

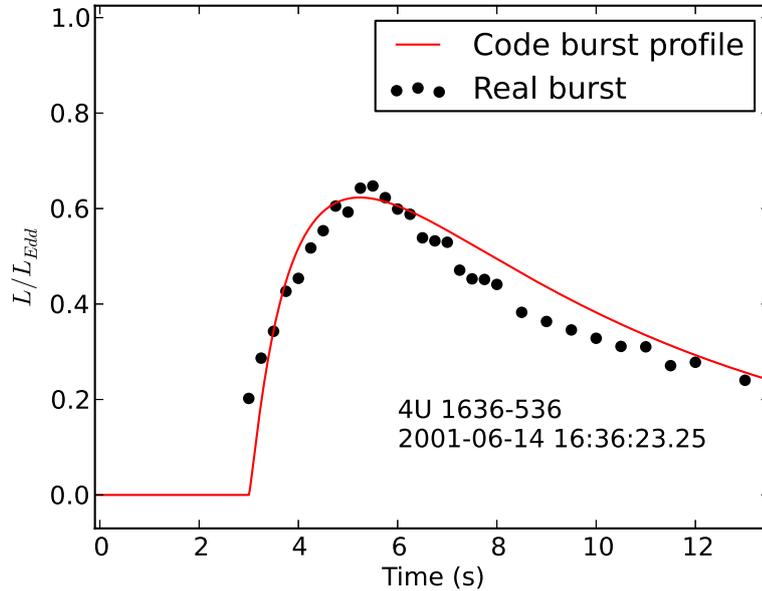


Figure 6.4 The simulated burst profile compared with a typical type I burst from the source 4U 1636-536. I have empirically matched the real light curve with the product of a power law and a modified exponential (see equation 6.3.1).

The simulated burst profile reaches a maximum luminosity of 62.3% Eddington after 2.25 s.

6.4 Prescriptions for β

6.4.1 Time-independent β

The correct way to determine $\beta(x, y, z)$ for any SPH particle is to integrate over the line of sight from the neutron star surface to the point (x, y, z) . However, doing this at every time step is computationally demanding and mathematically difficult. I devised an alternative scheme by making a few simplifying assumptions. First, I assumed that the disc does not change structure appreciably during a burst. This allows me to prescribe $\beta(R, z)$ for cylindrical radius R and distance from the midplane z at the beginning of the simulation and hold it fixed for all time. I also assume light travelling directly from the star to any point inside the disc along a direct path is obscured by the disc, but that light can reach that point by diffusing to that point in a direction normal to the plane of the disc. That is, points inside the disc are primarily lit by starlight scattered down from above and not directly from the star. Thus the optical depth for vertical diffusion of a point in the disc is

$$\tau = \kappa \int_{\infty}^{|z|} \rho(R, z') dz'. \quad (6.4.1)$$

After a little algebra, I find

$$\tau = \sqrt{\frac{\pi}{2}} \kappa \Sigma(R) H(R) \left[1 - \operatorname{erf} \left(\frac{|z|}{\sqrt{2} H(R)} \right) \right], \quad (6.4.2)$$

and

$$\beta = L_* \exp(-\tau). \quad (6.4.3)$$

I ran a low-resolution simulation with this β prescription and 20,000 particles, and let it settle for 30,000 timesteps (1.0 s). I then illuminated it with $L = L_{\text{Edd}}$ for another 30,000 timesteps to test whether a high luminosity will cause large changes in the accretion disc structure. Figure 6.5 shows the distribution of SPH particles in the settled disc and the burst-modified one. Two prominent stripes of material above and below the disc are evident. These structures consist of material lifted above the disc into regions where β is fixed at 1 and continue to experience large radiation forces even though, in reality, they

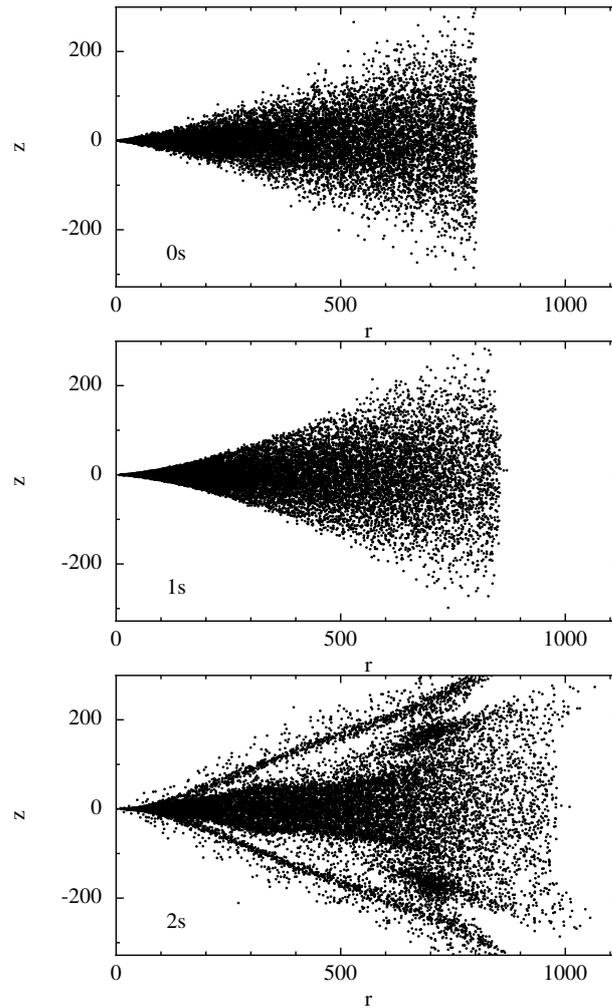


Figure 6.5 Test of the theoretical β prescription (Equations 6.4.2, 6.4.3) with 20,000 particles, demonstrating the formation of spurious structures resulting from fixed β . The top panel shows the initial configuration of the disc, the middle panel shows the disc after evolving without any luminosity for 1 s, and the bottom panel shows the disc after having been illuminated by $L = L_{\text{Edd}}$ for one second. The disc retains its initial shape in the absence of any radiation force, but its structure changes dramatically for this β prescription when strong radiation forces are present. Two prominent stripes of material are evident above and below the disc. The requirement that the disc not change structure significantly is not met, so this method of prescribing β is not suitable.

would partially self-shade and not experience such a great effect. The structures are not physically plausible, and the requirement that the disc not change structure significantly is clearly not met.

As the time-independent β prescription proved to be inadequate, I devised several new ones based on the actual distribution of SPH particles at each timestep.

6.4.2 Alternative approaches

Now I consider four alternative approaches to approximate the transport of radiation through the accretion disc. I illustrate them in Figures 6.7, 6.8, and 6.9 using a disc setup with 100,000 particles and illuminated with $L = L_{\text{Edd}}$.

Approach A

We do not allow for absorption in the disc at all, and simply allow the entire disc to be illuminated by all the luminosity of the star. In this approach the opacity of the disc material is zero.

Approach B

In this approach I integrate and absorb the luminosity along the line of sight from the neutron star, which I assume to be an emitting point particle.

I retained the assumption symmetry about the midplane, and divided the space into a spherical polar grid with $(n_r, n_\theta, n_\phi) = (63, 16, 24)$ divisions in each dimension. These numbers give a balance between the need for good resolution and the need to populate the bins with a sufficient number of particles. The radial divisions range from the surface of the neutron star at 10 km, to the initial edge of the accretion disc at 8000 km. They are unevenly spaced with the inner radius of the i th grid point given by

$$r_{\text{in}}(i) = 10\text{km} \times Ai^2, \quad (6.4.4)$$

where

$$A = \left(\frac{8000 \text{ km}}{10\text{km}} \right) \frac{1}{n_r^2}. \quad (6.4.5)$$

The θ angular divisions are also unevenly spaced and range from zero to $\pi/2$, with the inner angle of the i th grid point being given by

$$\theta_{\text{in}}(i) = \frac{\pi\sqrt{i}}{2\sqrt{n_\theta}}. \quad (6.4.6)$$

This uneven grid spacing is adopted to ensure good resolution near the neutron star and near the midplane, where most of the particles are. The ϕ grid bins are evenly spaced.

From the (x, y, z) coordinates of every particle I calculated its position in the grid:

$$\begin{aligned} r &= \sqrt{x^2 + y^2 + z^2}, \\ \theta &= \arccos(|z|/r), \\ \phi &= \pi + \text{atan2}(y, x) \end{aligned} \tag{6.4.7}$$

where θ is the angle measured from the north pole, $\phi \in [0, 2\pi)$ is the azimuthal angle, and atan2 is the usual two-argument inverse tangent function for returning the correct quadrant of the calculated angle. I then calculate the density in that cell. The volume of a cell with inner and outer radii r_{in} and r_{out} , and limiting angles θ_{min} and θ_{max} is

$$V_{\text{cell}} = \frac{2\pi}{3} (r_{\text{out}}^3 - r_{\text{in}}^3) [\cos(\theta_{\text{min}}) - \cos(\theta_{\text{max}})]. \tag{6.4.8}$$

I calculate the optical depth τ for each cell by numerically integrating from the neutron star surface to the cell midpoint:

$$\tau(r_{\text{mid}}, \theta_{\text{mid}}) = \int_{10\text{km}}^{r_{\text{mid}}} \kappa \rho(r', \theta_{\text{mid}}) dr' \tag{6.4.9}$$

using the trapezoidal rule, where $r_{\text{mid}} = \frac{1}{2}(r_{\text{in}} + r_{\text{out}})$ and similar for θ_{mid} . Finally,

$$\beta(r_{\text{mid}}, \theta_{\text{mid}}) = \frac{L_*}{L_{\text{Edd}}} \exp[-\tau(r_{\text{mid}}, \theta_{\text{mid}})]. \tag{6.4.10}$$

The parameter β can then be calculated for any particle by bilinear interpolation on the grid. I actually interpolate on the logarithm of β to avoid errors involving steep β gradients, which could result in calculated values becoming negative. The interpolation increases the resolution of β , making up for the large bin sizes.

$$\begin{aligned} (x_{i+1} - x_i)(y_{i+1} - y_j)f(x, y) &\approx (x_{i+1} - x)(y_{j+1} - y)f(x_i, y_j) \\ &+ (x_{i+1} - x)(y - y_j)f(x_i, y_{j+1}) \\ &+ (x - x_i)(y_{j+1} - y)f(x_{i+1}, y_j) \\ &+ (x - x_i)(y - y_j)f(x_{i+1}, y_{j+1}), \end{aligned} \tag{6.4.11}$$

where i, j are indices of gridpoints and $x_i < x < x_{i+1}$ and $y_j < y < y_{j+1}$.

Approach C

W92 argued that the disc should develop turbulent instabilities because angular momentum losses in the surface layers of the disc cause the Rayleigh criterion (specific angular momentum increases with distance from the midplane) to no longer be satisfied. In that work, these instabilities were assumed to efficiently mix the surface with the interior, thus allowing angular momentum loss to be distributed evenly in the vertical direction. W92 therefore felt justified in neglecting the vertical structure of the disc, simplifying the 1D simulations considerably.

I point out that such instabilities should be self-suppressing. Subtracting an equal amount of specific angular momentum from every particle in some vertical distribution causes the Rayleigh criterion to be satisfied again, thus removing the driver of the turbulence. This phenomenon is therefore likely to produce only partial mixing. The magneto-rotational instability, which was not widely known at the time of W92, may produce turbulence of a similar kind. It may, in fact, be a dominant mechanism for angular momentum transfer inside discs (Balbus & Hawley 1998). Furthermore, it is not obvious that turbulent timescales are negligible compared to the dynamical timescale near the NS surface. I therefore conclude that vertical structure cannot be neglected when I treat angular momentum losses. Nevertheless, this β prescription accounts for the possibility that angular momentum can be lost at sites distant from where the radiation is absorbed.

In this approach I calculate β for each particle as in approach B. I then divide the space into a one-dimensional cylindrical grid (i.e., I also abandon the ϕ divisions), and assign to every particle within a cylindrical division the average β of the particles in that division. This approximates W92's approach of distributing the radiation force impinging on an annular disc segment evenly and instantaneously throughout its vertical structure. The cylindrical radial divisions have the same radii as the spherical divisions.

Approach D

As discussed in 6.4.2.3, vertical mixing may only be partial. In that case, most of the angular momentum loss will occur in the outer layers of the disc, with regions near the midplane less affected. The same thing will happen if light impinging on the disc is partially

scattered downward into it. I have therefore devised a radiation transport approximation where β is found by integrating density in the z direction. This is similar to the first attempt at prescribing β in §6.4.

I bin the particles into a cylindrical grid, with $(n_r, n_\phi, n_z) = (63, 16, 30)$ and densities in each cell determined by particle counts and cell volumes. I have

$$\tau = \int_{-\infty}^0 \rho \kappa dz \quad (6.4.12)$$

for cells below the midplane and

$$\tau = \int_{\infty}^0 \rho \kappa dz \quad (6.4.13)$$

for cells above the midplane. Again, $\beta = \exp[-\tau]$. This approach communicates β to the midplane but decreases the effect on the surface layers, as shown in Figure 6.8. Because the illuminated surface area increases with increasing radius, and the number of particles decreases, β also increases with increasing radius. It is also sensitive to particle number, and the position of particles near the surface layers (as these contribute the most to β), so $\beta(R)$ is not cleanly monotonic.

6.4.3 Feedback Luminosity

Material accreting onto the star loses gravitational potential energy, which is released as radiation (Equation 1.1.2):

$$L_{\text{Acc}} = \frac{GM_* \dot{M}}{R_*}. \quad (6.4.14)$$

Since it is emitted from the stellar surface it illuminates the disc, contributing to the radiation drag. Thus the total luminosity of the star can be expressed as

$$L_* = L_{\text{Burst}} + E_F L_{\text{Acc}}, \quad (6.4.15)$$

the sum of luminosity from nuclear burning, which I insert directly into the code as a time-variable prescription, and feedback luminosity, which depends on the instantaneous accretion rate. I have introduced the feedback luminosity efficiency factor E_F as a means of gradually introducing the feedback luminosity. In the simulations I increase it linearly from zero at 1 s before the burst to unity at the onset of the burst.

For discs such as this one, where the accretion rate is low, it turns out to make little difference whether feedback luminosity is included or not. I verified this by running two low-resolution simulations with 80,000 particles, one with feedback luminosity brought in at $t = -1$ s and one with no feedback luminosity. The difference between the two is only about 0.06%, which is very small compared to the angular momentum losses caused by burst luminosity, as discussed in §6.5.

In high accretion rate systems, such as Cyg X-2, where the neutron star can accrete at almost the Eddington limit, radiation drag caused by the luminosity from accretion is likely to be a very important effect.

Accretion of individual SPH particles

As mass in SPH is represented as discrete particles, the accretion luminosity can only take on values that are a fixed multiple of GM_*M_{part}/R_* , where M_{part} is the mass of one SPH particle. This can potentially lead to “flickering” of the star. To determine whether this can affect particle trajectories I ran the test code again as above, but at every timestep I let the luminosity be either zero or $L = 0.2L_{\text{Edd}}$ with equal probability (for an average of 10% Eddington). The two trajectories are indistinguishable. I am therefore confident a large number of rapid and discrete accretion events over a time interval will behave the same as steady accretion with the same average luminosity. Thus, there is no danger in allowing the star to flicker as individual SPH particles are accreted. For low pre-burst accretion rates the addition of feedback luminosity is negligible anyway.

6.5 Results

6.5.1 Full simulations

In the following pages I show plots of the f_a evolution of the disc for approaches A, B, and D (§6.4). Figure 6.11 shows the results of the full simulations with 2.5 million particles. I calculated f_a for the simulations as follows:

I calculated accretion rate by counting every particle accreted over 2,500 timesteps (0.083 s, a third of the size of the time bins of *RXTE*). I then scaled to f_a by defining the last 2,500 timesteps before the beginning of the burst to be $f_a = 1$.

It is evident that approaches A and D, in which radiation drag affects interior regions of the disc, show a strong increase in accretion rate. These approaches reach peak f_a values for 7.87 for approach A and 1.75 for approach D. For approach A, the rise in f_a is faster than the increase in burst luminosity while in approach D the temporal behaviour in f_a follows that of the burst luminosity closely. In approach B there is almost no response in f_a , suggesting that angular momentum loss is inefficient for this prescription. It only attains a peak f_a of 1.035, i.e., a 3.5% increase in accretion rate.

Due to a coding error that was discovered late and the very long run times of the simulations, there has not been time to perform a full simulation of approach C. Approach D has not had time to finish running, but the most interesting times have already been simulated.

In the following pages I show the behaviour of total disc angular momentum. As before, the biggest effect is for approaches A and D. These lose 16.3% and about 14% of their initial angular momenta over the course of the run, as shown in Figure 6.18. Approach B shows a decrease of 5.5% of its initial value over the entire run, but it has not resulted in a significant enhancement of accretion rate.

6.5.2 Low resolution simulations

I have also run the four radiation prescriptions for 250,000 particles. As discussed in §6.2.2, this is not sufficient to resolve the scale height of the disc near the star, but the run times are short enough that approach C can be included. The accretion enhancements are shown in Figure 6.19. By comparing these with Figure 6.11 it is evident that the lack of resolution of scale height does not greatly affect the qualitative behaviour. Approach A reaches a slightly lower peak f_a of 6.97, compared with 7.87 for the full simulations, and the very slight increase for approach B is somewhat more prominent, reaching a peak f_a of 1.25. The behaviour of Approach C is similar, attaining peak f_a of 1.22.

The angular momentum losses are shown in Figure 6.18. Approach A behaves almost identically as for the high resolution run, losing 16.5% of its initial value. Approach B loses

somewhat more angular momentum, 8.9%, than it did in the high resolution run. However, it does not result in a commensurate increase in f_a .

Because approaches A, B, and D behave the same for both high and low resolution I expect that approach C would have behaved the same for 2.5 million particles as it does for 250,000.

6.6 Discussion

The results of §6.5 show that significant accretion rate increases due to radiation drag are possible for several physically plausible approaches to radiation transmission. I have found that accretion enhancements of $f_a \approx 8$ can be achieved for a disc that is entirely illuminated (A), and $f_a \approx 2 - 3$ for discs that partly self-shade but in which angular momentum loss is transferred to the midplane (C and D). Only discs in which the midplane is entirely shaded show no increase in accretion rate, though they do still lose a significant part of their angular momentum (see Figures 6.12 and 6.18).

As shown in Figures 6.14 and 6.17 it is possible for a burst to simultaneously induce both inward and outward motions in the disc. These Figures show that the inner disc has been significantly depleted compared to a relaxed, unilluminated disc (Figure 6.13), and also exhibit an expanding torus-like structure. The expanding structure is caused by outward radiation pressure.

Since the simulated discs were accreting at 2.5% Eddington prior to the burst, an accretion enhancement of a factor of a few does not reach the 37% Eddington observational upper limit for non-PRE bursts found in 5.3, but enhancements of up to about 20% appear plausible. I have been unable to reproduce the extremely high f_a values of W92.

As shown in Figures 6.11 and 6.19, the behaviour of Approach D most closely resembles that of the measured f_a for the selected burst. However other bursts, even with similar light curves, can exhibit different f_a behaviour. In Figure 6.20 I show burst light curves for a selection of six non-PRE bursts, showing that f_a behavior can differ dramatically between bursts. Some, such as the burst from KS 1731–260, reach large f_a rapidly, qualitatively similar to Approach A. Others, such as the burst from 4U 1636–536, show a more gradual

and longer-lasting increase resembling Approach D. Bursts can also behave differently to any of the four approaches in this work; for example the burst from 4U 1702–429, which reaches peak f_a significantly later than the peak of the light curve.

The discs I have run are locally isothermal, with temperature specified as a function of radius. Heating of the disc by radiation is not modelled. I expect that heating the disc would amplify the accretion rate increase, primarily because this increases the scale height of the disc. Clearly, a disc that extends further vertically will absorb more radiation from the star. The density of the disc will also decrease, as will the temperature-dependent opacity of the disc material (Lewin et al. 1993 eq. 4.13b), thus lowering the optical depth and allowing radiation to penetrate further into the disc. As I have found, the degree to which radiation can get in to the disc is a crucial factor in determining the \dot{M} response. Even in the absence of radiation effects, discs with a larger scale height accrete more rapidly ($\dot{M} \propto H^{20/3}$ for a Shakura-Sunyaev disc, for example; Frank et al. 2002).

I have not taken into account the fact that the star is an extended emitter. For a geometrically thin disc that extends to near the stellar surface, some of the emission will strike disc from above the midplane. In that case approaches B and C might have some of the character of approach D, and may accrete faster.

The disc initially extends to the neutron star surface. For low pre-burst accretion rate, this is an unrealistic assumption. It is likely that the inner edge of the disc is truncated at some radius above the neutron star surface, with material passing this point accreting as a fast-moving optically thin wind. It would perhaps have been better to set the accretion radius to something significantly larger than the neutron star radius because then the unresolved scale height of the inner disc lower resolution simulations (§6.5.2) may in fact approximate the true physical behaviour than the full simulations.

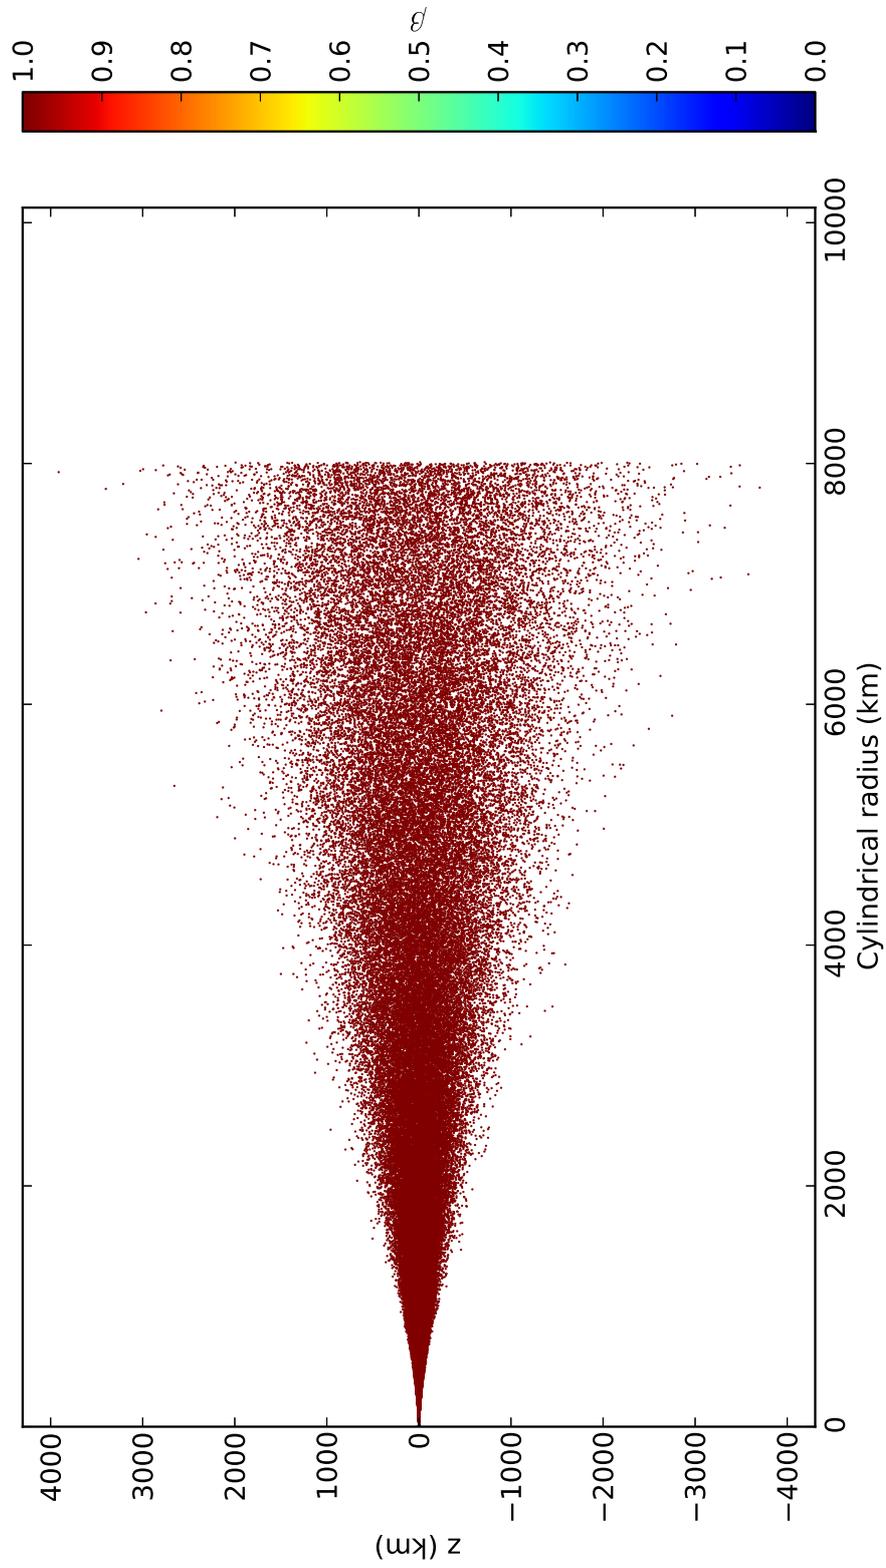


Figure 6.6 Distribution of SPH particles for approach A, with the β applied to them indicated by their colours. In this approach all the particles are fully illuminated.

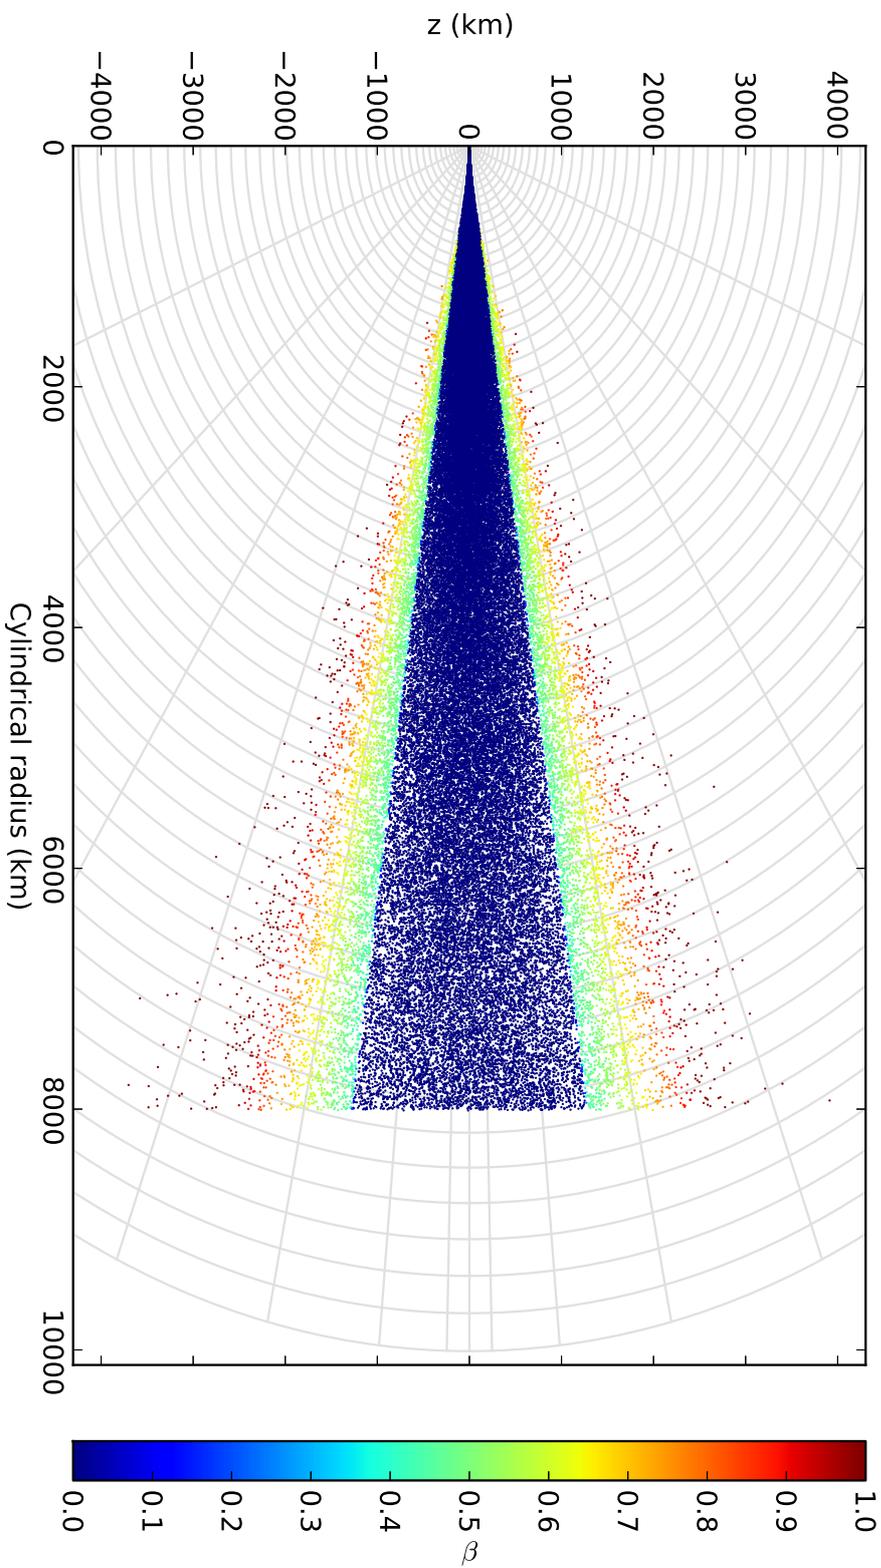


Figure 6.7 Distribution of SPH for approach B, with the β applied to them indicated by their colours. In this approach absorption towards any point inside the disc is treated by integrating from the origin to that point and interpolating on a spherical grid (radial and polar divisions shown as grey lines; azimuthal divisions not shown). It is clear that most of the radiation forces affect the parts of the disc far from the midplane, with the interior effectively shaded.

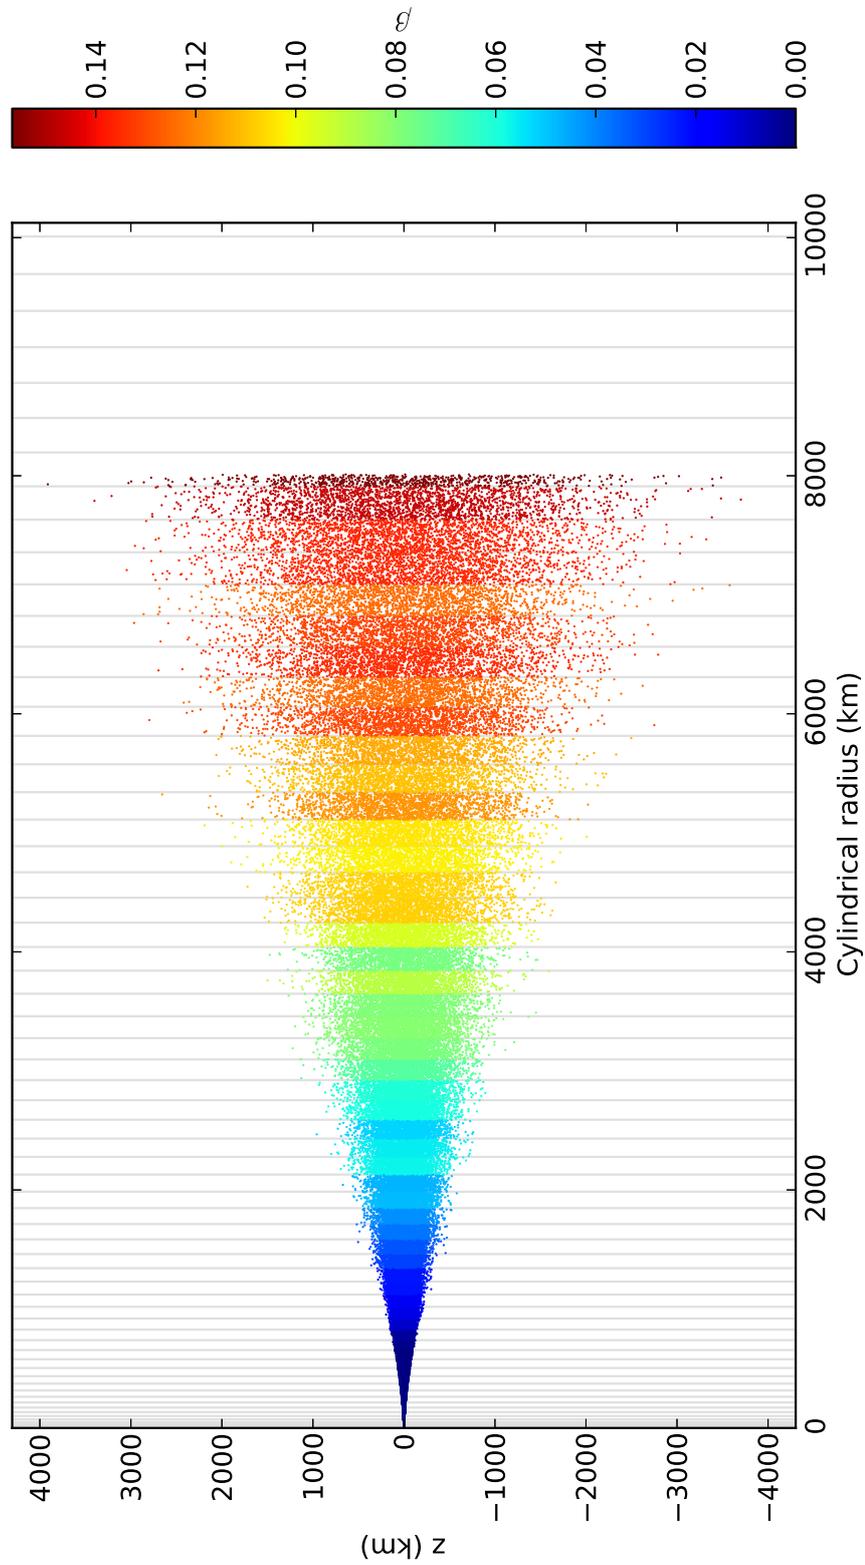


Figure 6.8 Distribution of particles for approach C, with β indicated by their colours. This approach is similar to approach B, except that I now average the β applied to the particles among all the particles in an annular ring (grey lines). The averaging communicates radiation drag to the midplane, but decreases the effect on surface particles; note that the maximum β for this description is about 0.16, whereas for the others it is 1.

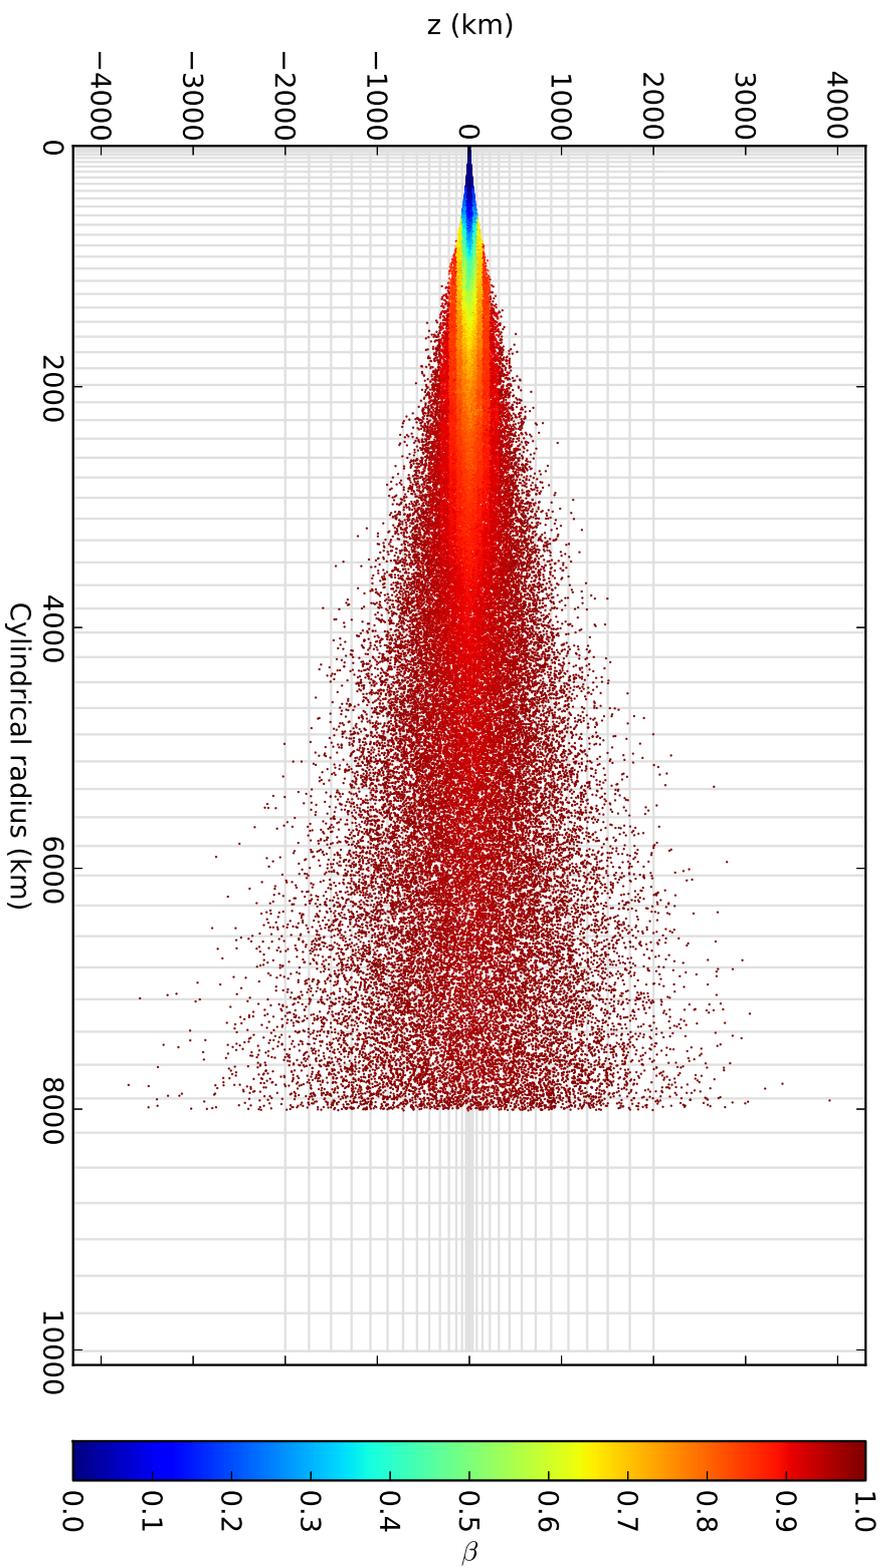


Figure 6.9 Distribution of particles for approach D, with the β applied to them indicated by their colours. Here I have attempted to approximate the effects of partial mixing through turbulence and scattering of light towards the midplane by integrating rays of light vertically on a cylindrical polar grid (radial and vertical divisions shown as grey lines). In this approach the outer layers of the disc are still illuminated strongly, but some light also reaches the midplane.

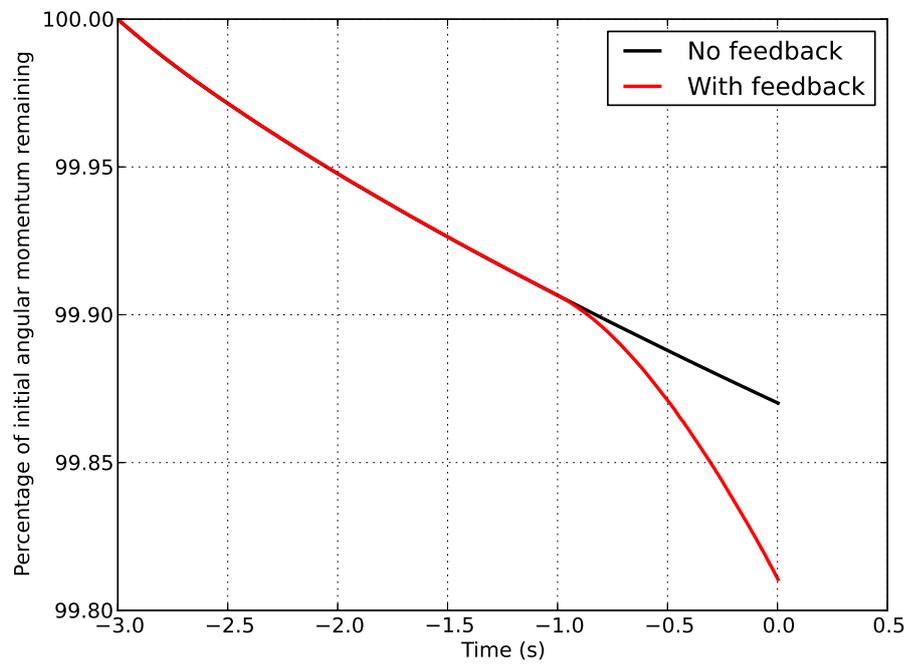


Figure 6.10 Angular momentum loss of a disc experiencing feedback luminosity (red curve) compared with one that does not (black curve). The efficiency factor E_F is increased linearly from zero at $t = -1$ s to unity at $t = 0$ s. The difference between the two is only about 0.06% of the initial angular momentum of the disc indicating that, for low accretion rate, the radiation drag caused by accretion luminosity has a negligible effect on the disc.

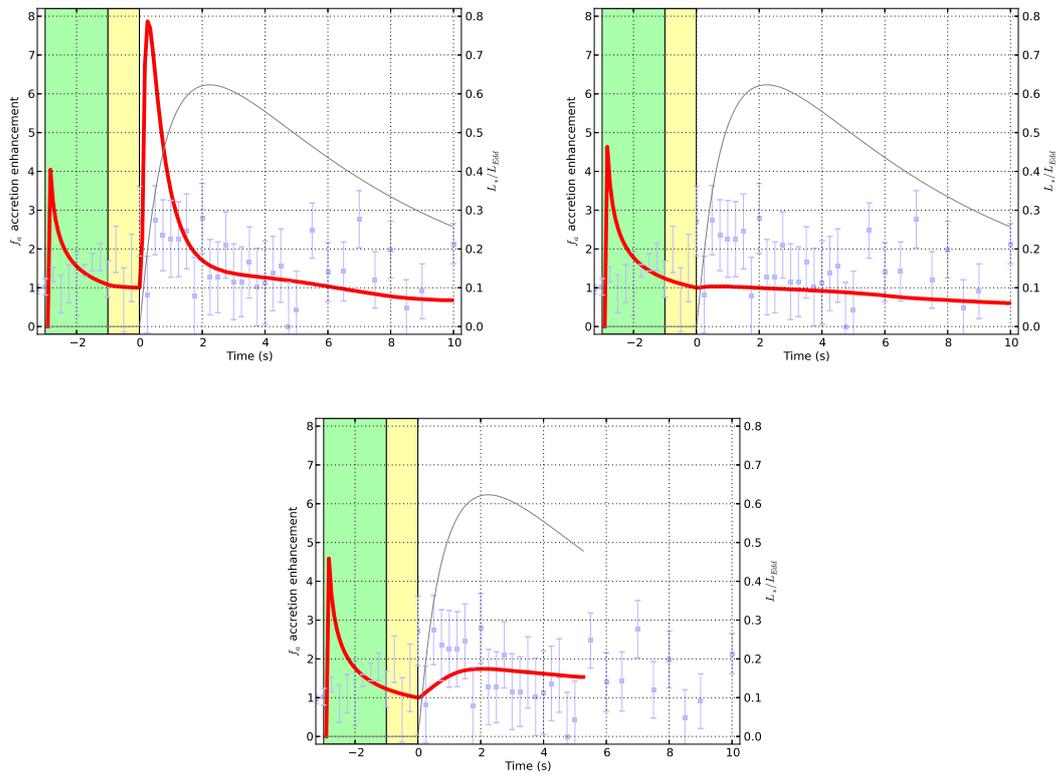


Figure 6.11 Accretion rate enhancement of simulations for 2.5 million particles for approaches A, B, and D. Approaches A and D show a large increase in accretion rate, but approach B shows almost no effect. Also shown is the adopted burst profile (grey curve) and the measured f_a for the model burst (grey points). The period of zero luminosity for removing the initial transient is shown in green, and the gradual introduction of feedback luminosity is shaded yellow.

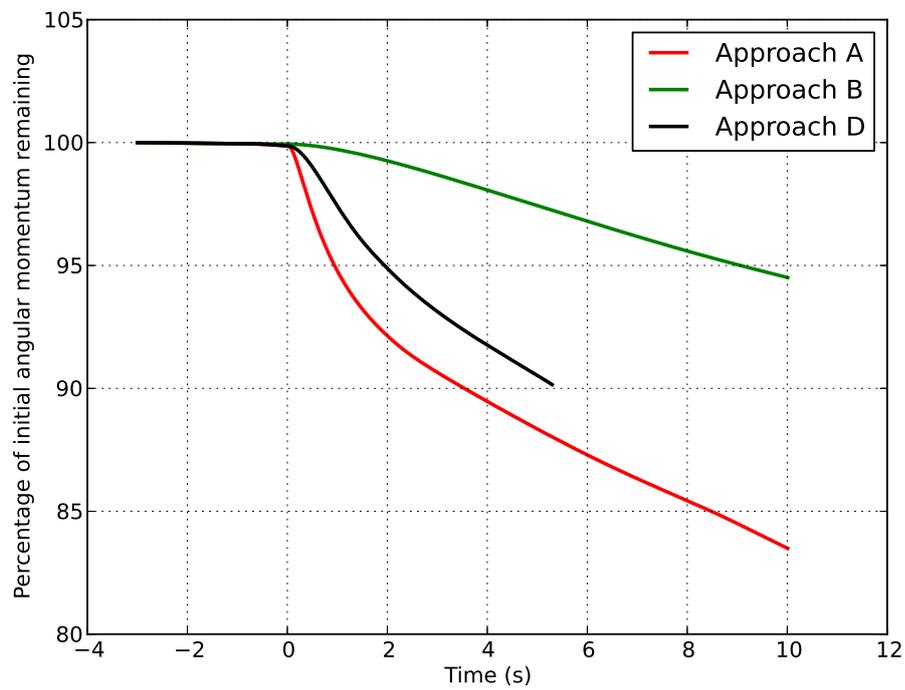


Figure 6.12 Angular momentum loss for approaches A, B, and D in the 2.5 million particle runs. Approaches A and D lose more than 10% of their initial angular momenta over the run, and approach B loses 5.5%. As expected, the approaches in which a large β is applied to the interior parts of the disc show more angular momentum loss than approaches where radiation drag affects only the outer parts of the disc.

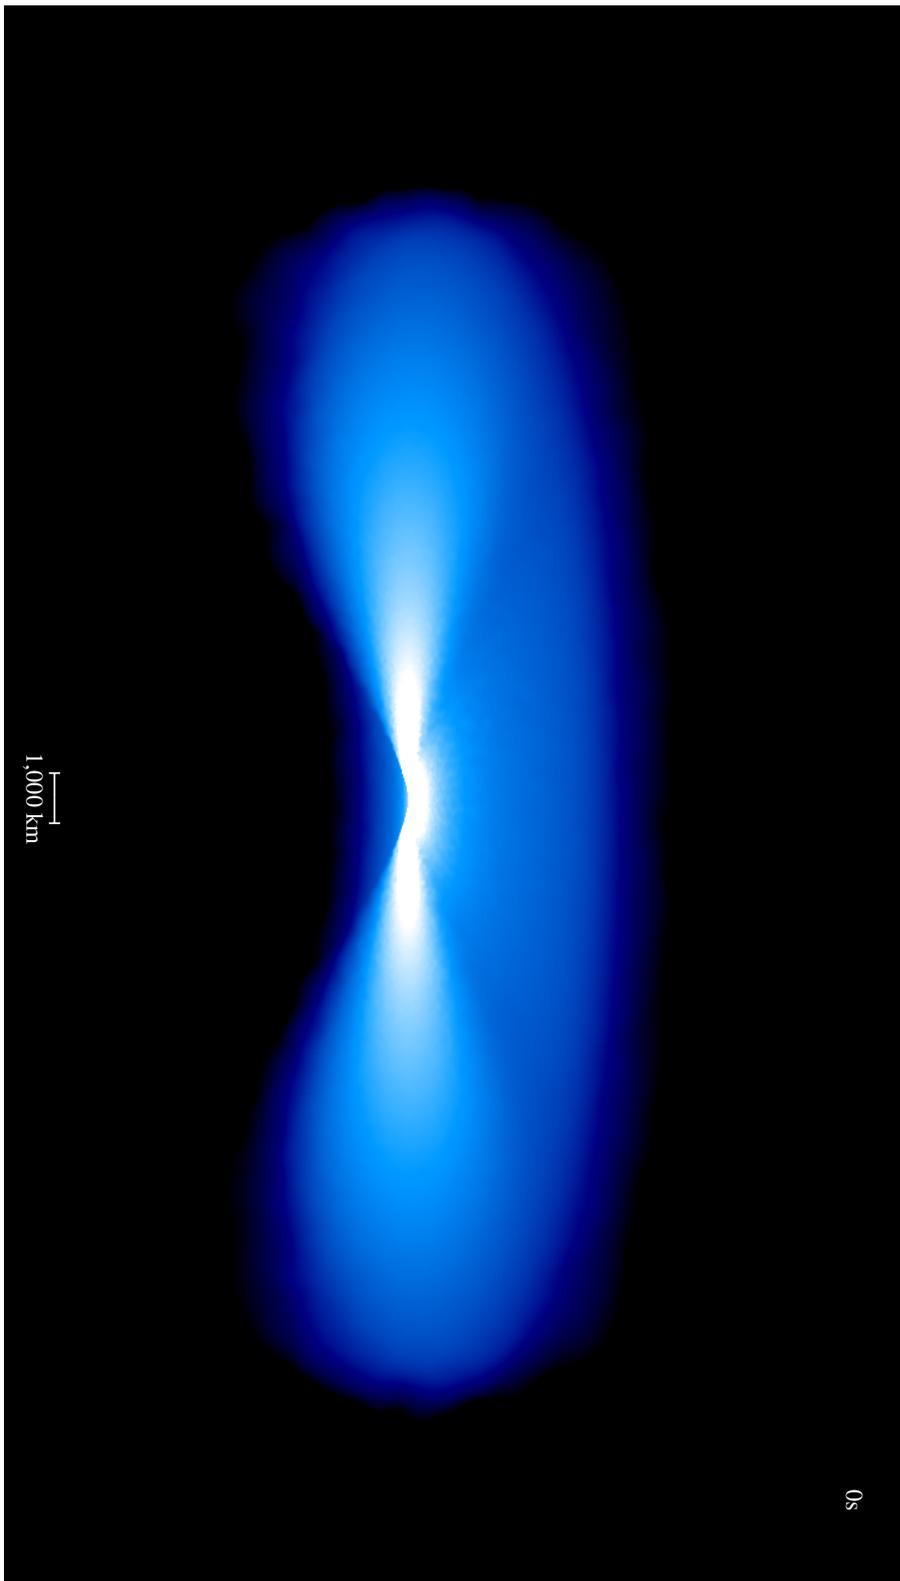


Figure 6.13 Rendered plot of the disc structure for the disc at 0 s, just before the onset of the burst, to show the disc structure before its modification by burst luminosity. The disc is shown tilted, with particles on the near side of the origin removed. Colors represent column density along the line of sight, with white indicating high column density. On this scale the neutron star itself is too small to be seen.

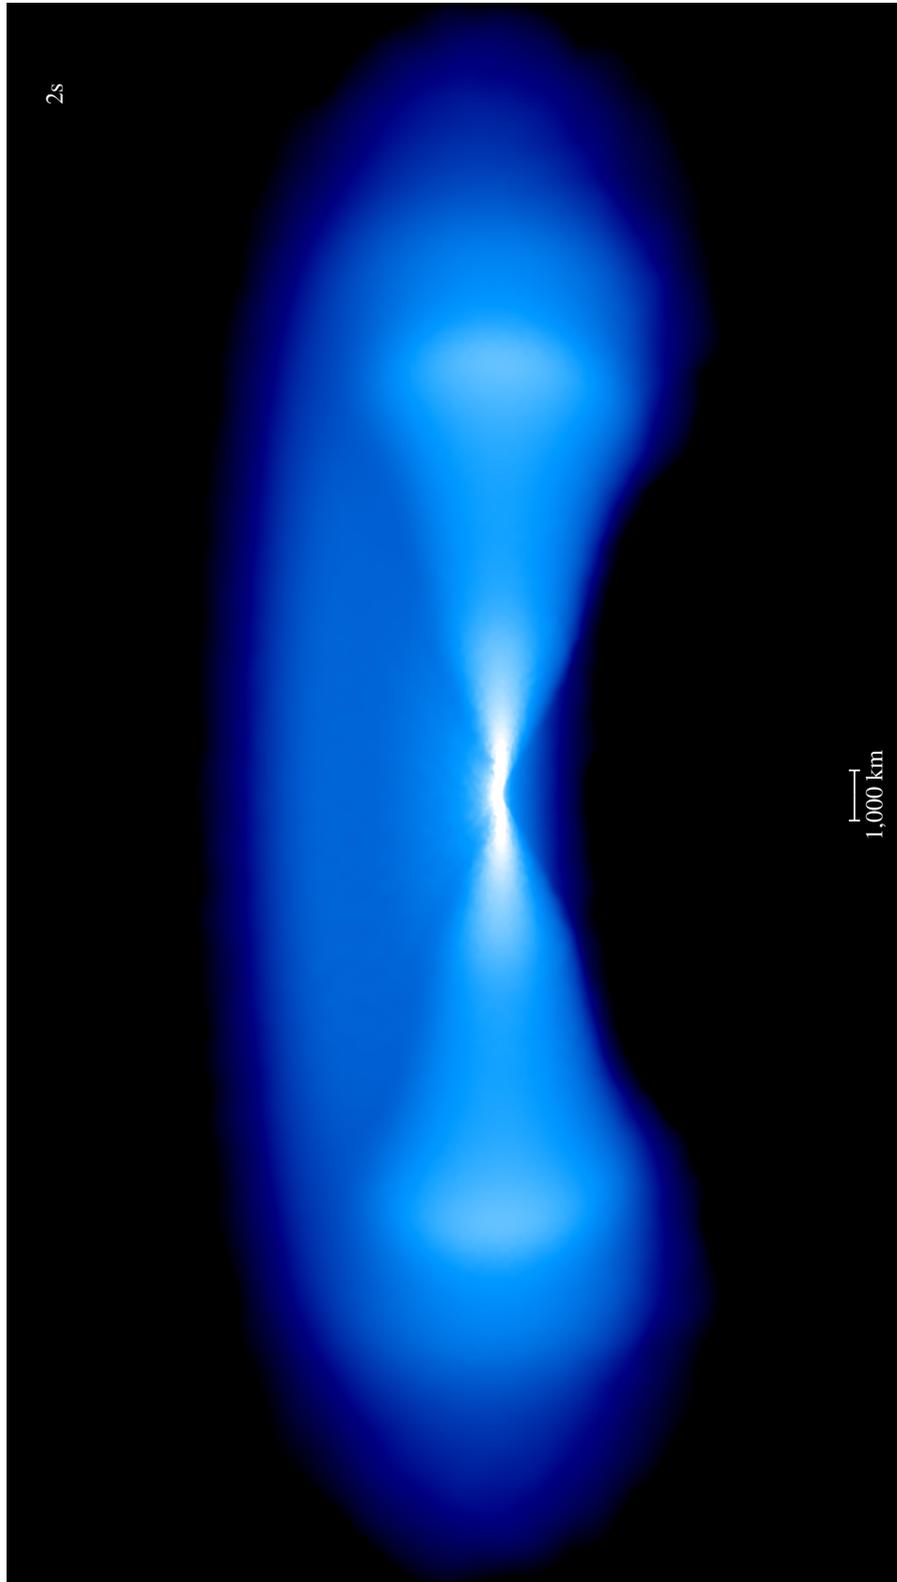


Figure 6.14 Rendered plot of the disc structure for Approach A after 2 s, showing considerable change in the disc structure. Compared with Figure 6.13 the inner regions of the disc have lost a considerable amount of material. Also visible is an expanding torus-like structure at $r \approx 600$ km caused by outward radiation pressure.

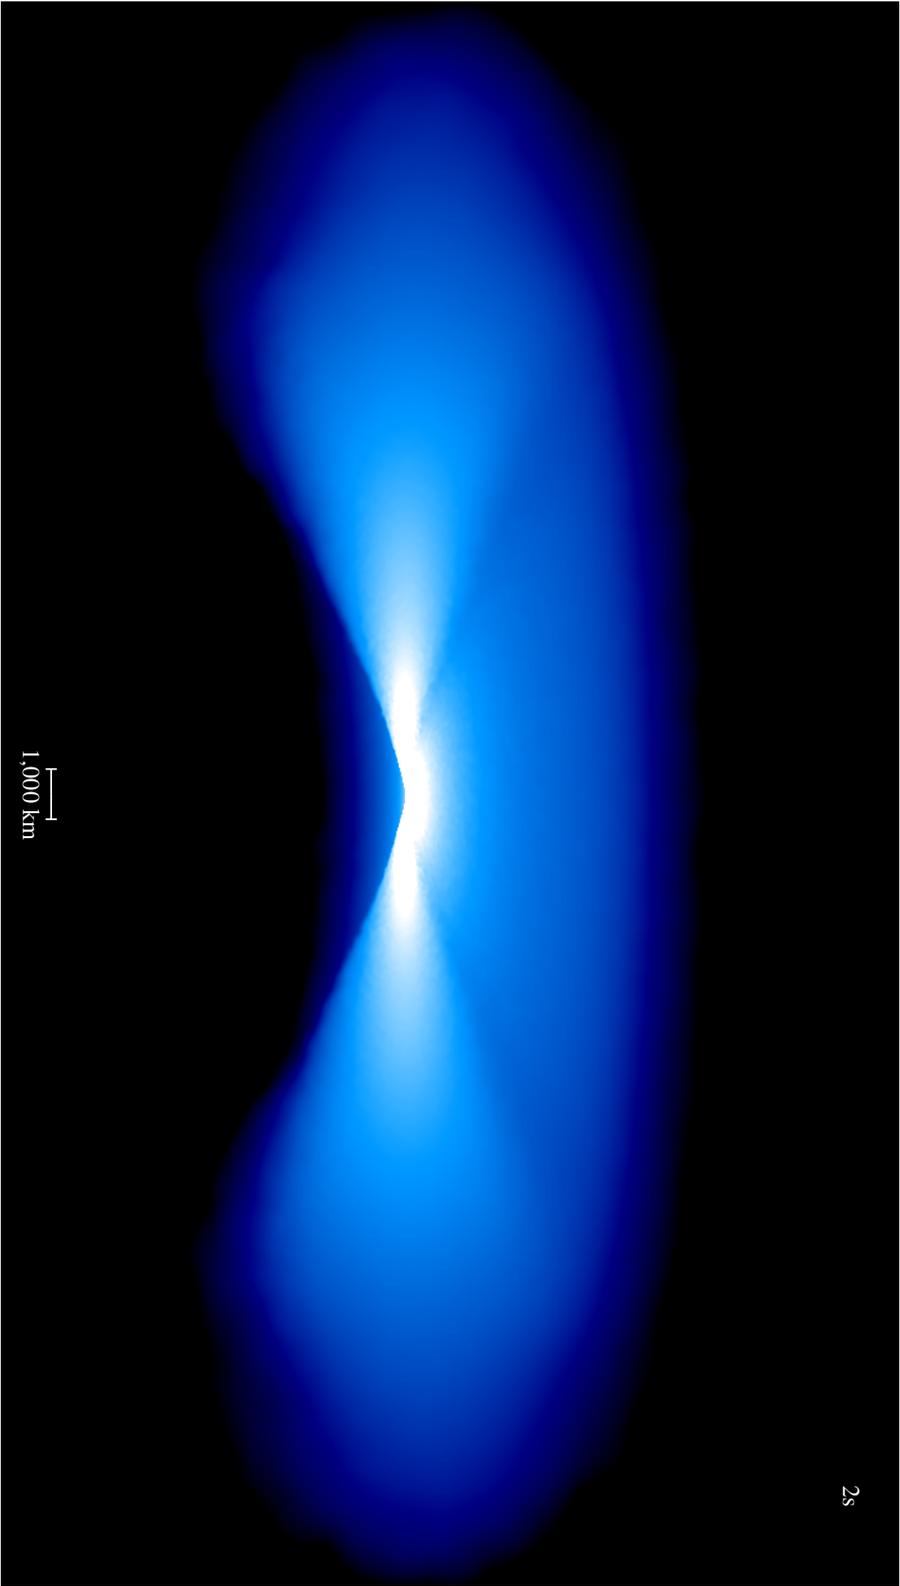


Figure 6.15 Rendered plot of the disc structure for Approach B after 2 s. The structure of the disc has hardly changed at all.

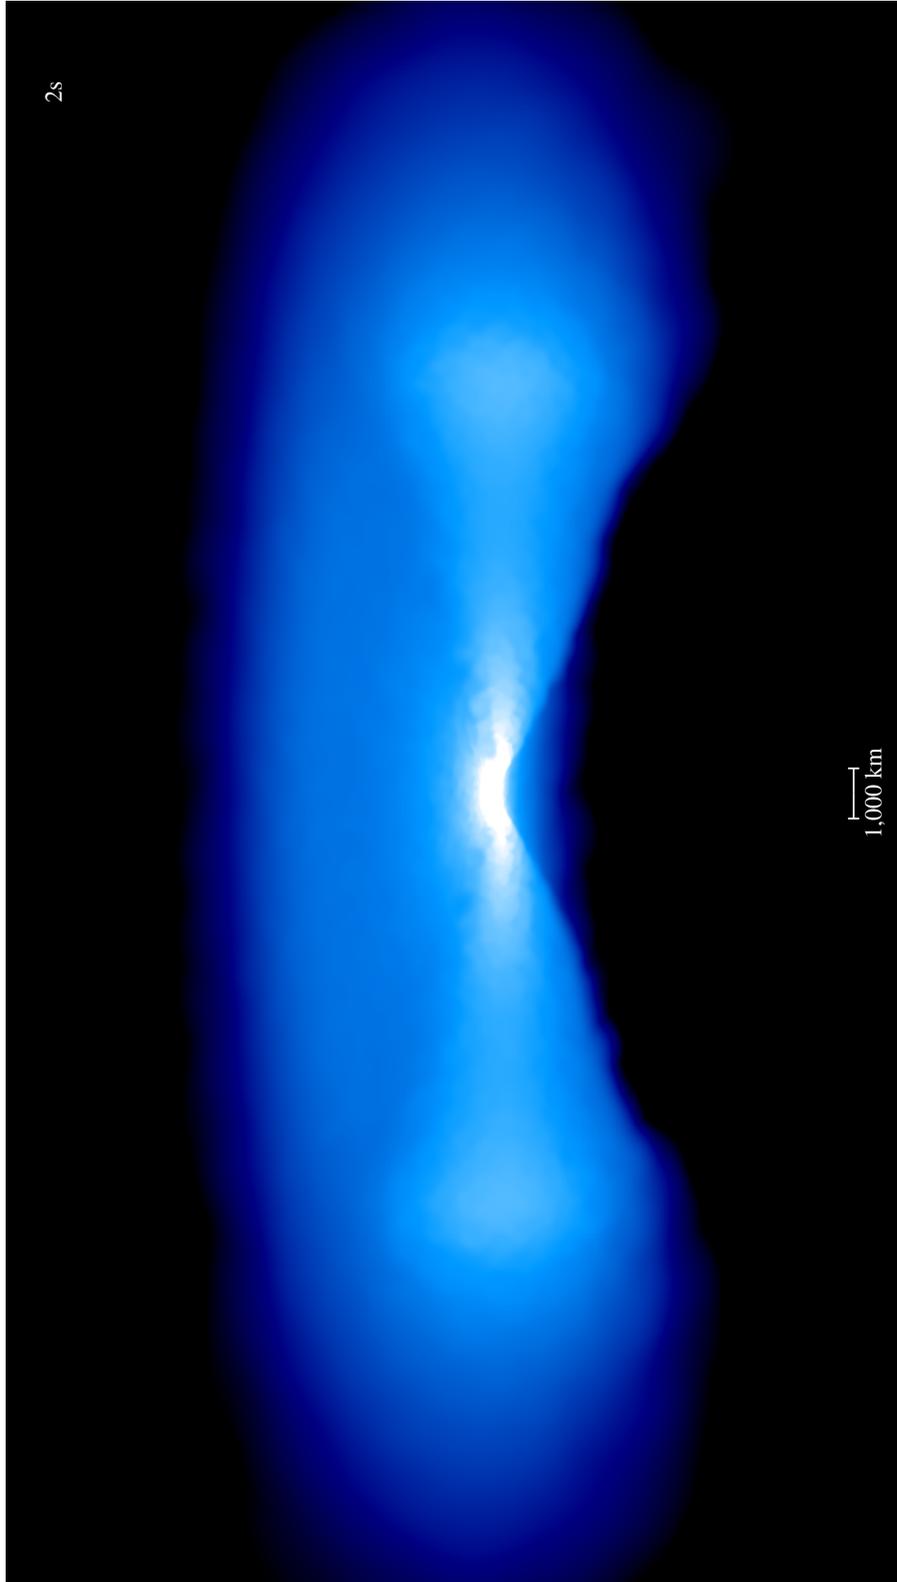


Figure 6.16 Rendered plot of the disc structure for Approach C after 2 s. Note that this Figure is of the 250,000 run, whereas the other rendered plots are all 2.5 million particles. The structure of the disc is virtually unchanged from the unilluminated disc shown in Figure 6.13.

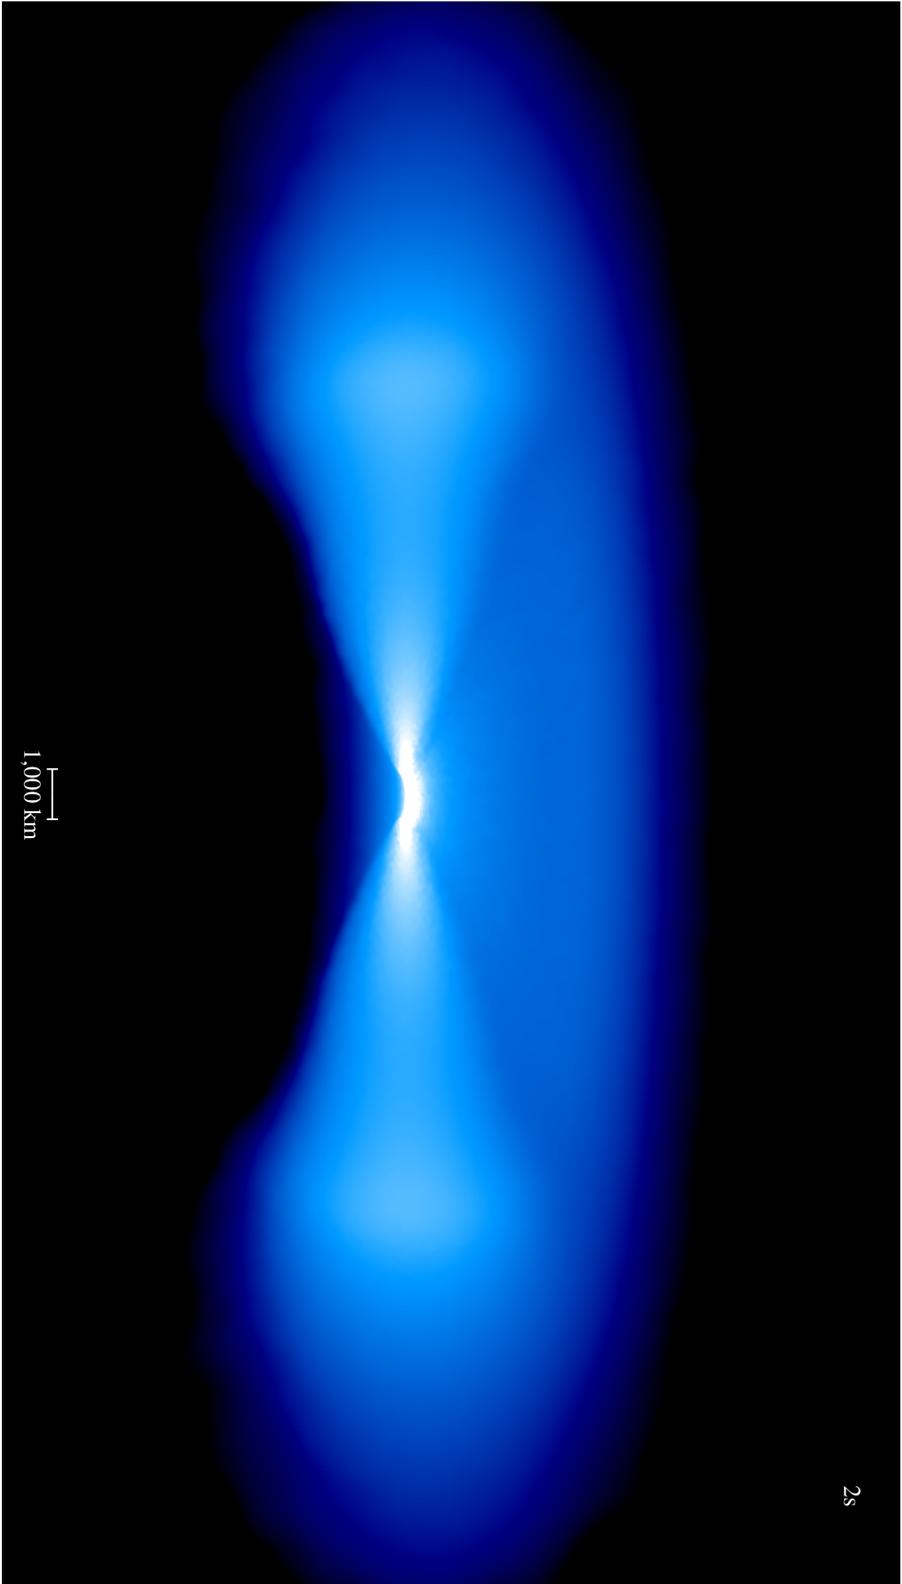


Figure 6.17 Rendered plot of the disc structure for Approach D after 2 s. It exhibits similar behaviour to 6.14, with an evacuation of the inner part of the disc through radiation drag, and a toroidal structure from radiation pressure.

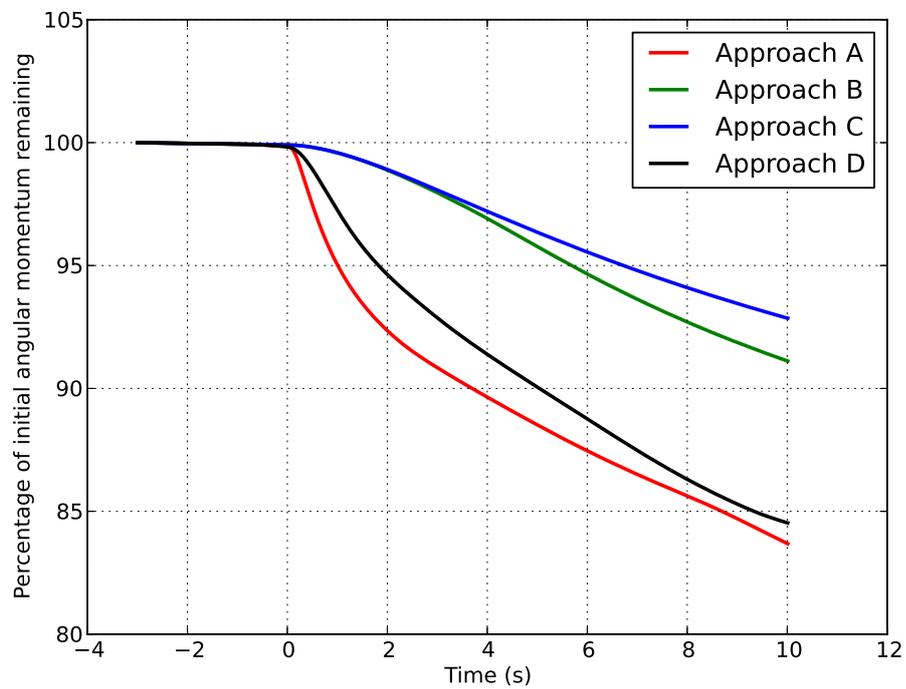


Figure 6.18 Angular momentum loss for approaches A, B, C, and D in the 250,000 particle runs. As with the full resolution runs, approaches A and D lose a substantial amount of their initial angular momentum. Approaches B and C lose comparatively less.

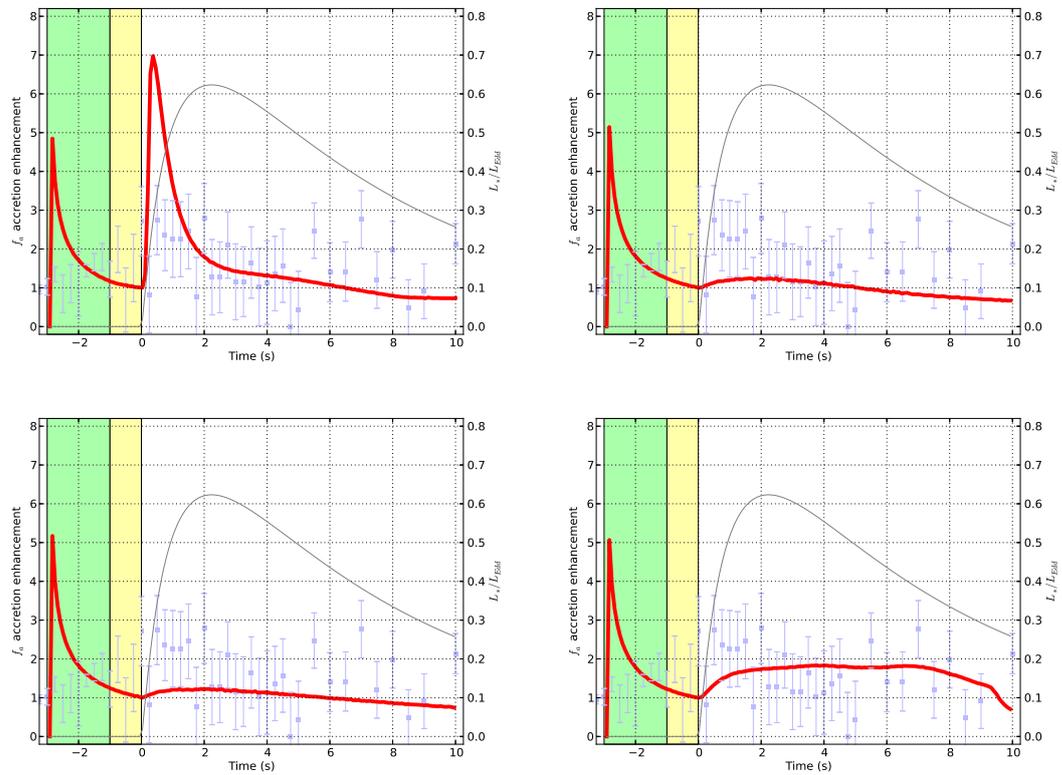


Figure 6.19 Accretion rate enhancement of the four radiation prescriptions for 250,000 particles. Also shown is the adopted burst profile (grey curve) and the measured f_a for the model burst. Although the scale height of the disc is not properly resolved for these simulations, the qualitative behaviour is not changed in approaches A, B, and D.

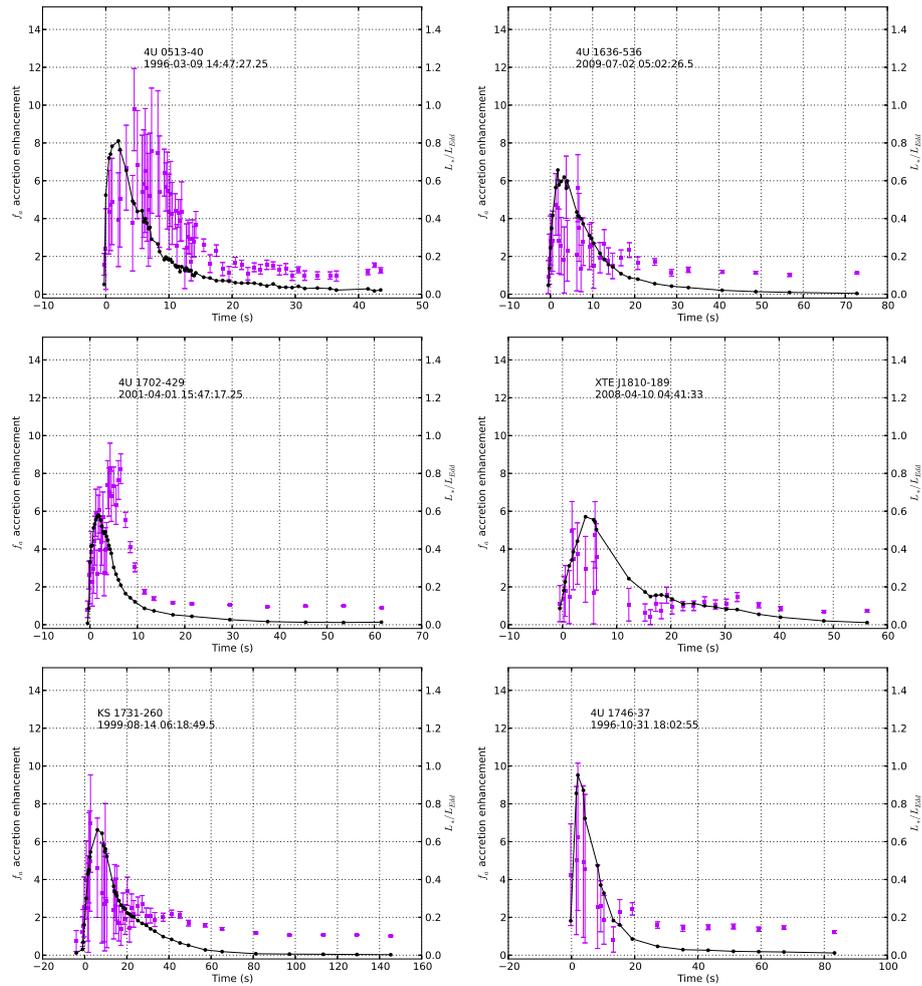


Figure 6.20 Plots of f_a (blue points) and total luminosity as a fraction of Eddington luminosity (black points) for six non-PRE bursts, from the analysis in Chapter 5, showing a variety of behaviour of f_a for similar burst profiles.

Chapter 7

Conclusions and Future Work

The first time you go through it, you don't understand it at all. The second time you go through it, you think you understand it, except for one or two small points. The third time you go through it, you know you don't understand it, but by that time you are so used to it, it doesn't bother you any more.

Arnold Sommerfeld

7.1 Discussion

I have performed observational and modelling investigations into whether or not the persistent emission contribution varies during type I X-ray bursts. I allowed the pre-burst emission to vary, parameterised by a factor which I denote f_a . I detected a statistically significant increase in f_a for nearly all the PRE bursts in the catalog, suggesting an enhanced (rather than suppressed) level of persistent emission during a burst. The new method results in a significant improvement in the reduced χ^2 of spectral fits compared to the standard model but not to the level of formal statistical consistency for all the spectra. Since the persistent emission is known to be an approximate measure of the accretion rate onto the neutron star, I interpret the results to indicate that the accretion rate onto the neutron star generally increases during bursts. Obviously the persistent spectrum could also change shape in response to a varying accretion rate, and there are suggestions ([Homan](#)

et al. 2007) that the X-ray luminosity may not be proportional to the mass accretion rate. There are, however, currently no predictions about how the persistent emission spectrum may change in response to rapidly increasing radiation drag, so I make the simplest assumptions that it does not change shape and that $f_a \propto \dot{M}$. I have shown that peak f_a measured during each burst anti-correlates with the pre-burst accretion rate with a slope consistent with that predicted by the theoretical models of Walker (1992), who investigated the effect of radiation drag on the accretion disc. If the effects of neutron star spin are accounted for, the magnitudes of the observed peak f_a are also consistent with the predictions of W92. This suggests that the detection of an increased pre-burst persistent emission reflects an enhanced accretion rate due to radiation torques on the accretion disc around the star.

All of W92's models are sub-Eddington, so some care must be taken in extrapolating their results to PRE bursts, so I have repeated the analysis for non-PRE bursts in Chapter 4. Even so, W92 points out that the expanding atmosphere in a PRE burst should have very little angular momentum compared with the accretion disc. Radiation coupling between disc and envelope can thereby spin down the disc, increasing the accretion rate. This would also have the effect of spinning up the envelope and boosting its expansion. Such a mechanism may help explain why the atmosphere takes longer to return to the neutron star surface than to reach maximum radius.

It may be argued that a deviation of the burst spectrum from a blackbody could mimic a variable accretion rate by introducing a high energy excess that the variable persistent flux method then attempts to remove. If the deviation is of fixed size then any spurious f_a measurement it causes will also anti-correlate with the pre-burst accretion rate γ , similarly to Figure 4.6. However, I would then also expect f_a to be constant with constant burst flux, but I found it to vary greatly during the Eddington-limited phase when the burst flux is approximately constant. I found no anti-correlation between f_a and blackbody temperature, which would occur if the blackbody component drops out of *RXTE*'s detection band and leaving only the deviation to be fit out. I also attempted to model a hard tail deviation in the burst spectrum using a power law with fixed index and normalization tied to the burst flux. This did not improve the spectral fits for the entire collection of bursts, and did not cause the detection of f_a to be suppressed. While in general I do not expect to see discrete

spectral features in PRE bursts, in superexpansion bursts they can be visible (in't Zand & Weinberg 2010). I have found that in these events f_a consistently drops rather than rises, consistent with zero flux from the source. As suggested by in't Zand & Weinberg (2010) this may be due to the emission from the burst component dropping out of the band detectable by *RXTE* together with the persistent emission being obscured by the superexpanding shell, though it is also possible that the vigorously expanding envelope disrupts the accretion disc and thereby temporarily halts accretion. in't Zand, Keek & Cavecchi (2014) argue that the spectrum observed in the latter part of the superexpansion episode, when the source is becoming detectable again, is a combination of a greatly expanded nova-like shell and the (partially obscured) neutron star and inner accretion disc. Measurements of f_a during this time are therefore compromised.

Nevertheless, an enhanced f_a cannot be attributed to the presence of spectral features superimposed on the burst component continuum. While I have attempted to exclude confounding spectral effects intrinsic to the burst component as an explanation for f_a enhancements, I cannot rule out the possibility of other interpretations.

If the variable persistent flux method causes the inferred photospheric radius at touchdown to differ systematically from the standard fits, then this would have implications for determinations of the equation of state of neutron star matter, such as those of Özel et al. (2009), Güver et al. (2010), Steiner et al. (2010). I investigated this by taking the ratio of the touchdown radii as determined by the variable persistent flux and standard methods for all the PRE bursts in the sample and obtained a mean value of 0.97 ± 0.11 . This indicates that the inferred touchdown radius can differ by $\sim 10\%$, and that there is a slight trend towards lower neutron star radii using the variable persistent method. I have found that f_a frequently remains elevated after the end of the Eddington-limited phase; this may have implications for studies that use the cooling tail to constrain the neutron star parameters (e.g., Galloway & Lampe 2012).

7.1.1 Structure of the accretion disc

The details of the transfer of material from the inner edge of the accretion disc to the neutron star surface are still uncertain (e.g., Bildsten 1998). It is not known if it occurs

at a close inner boundary layer, or if there is some mechanism that regulates the infall. PRE bursts offer a means of settling this question: to be subject to radiation pressure the expanding photosphere must be optically thick and therefore must obscure the observer’s view of everything interior to it. The atmosphere expanding so far that it covers all of the emitting parts of the accretion disc would cause the persistent emission to be reprocessed inside the optically thick atmosphere and effectively become part of the burst emission. Even very modest radius expansion would hide the boundary layer due to this “shrouding”, and cause f_a to decrease, since the boundary layer is thought to be geometrically small and located at a radius $\sim 1.2R_*$ (e.g., Popham & Sunyaev 2001). Thus the finding of consistently high f_a during a burst argues against the existence of a thin boundary layer that remains near the surface of the star and dominates the persistent emission (see Figure 4.2), but is consistent with a boundary layer that becomes much wider in response to increased luminosity as suggested by Popham & Sunyaev (2001).

I have extended the variable persistent flux method developed in this work and in’t Zand et al. (2013) to all type I bursts observed by *RXTE*. The new method gives superior spectral fits to type I burst spectra, whether these are radius expansion bursts or not. The detailed conclusions are as follows:

- i) The variable persistent normalization approach developed in this thesis applies as well to non-PRE bursts as it does to radius expansion bursts. The quality of spectral fits generally improves, and the intensity of the persistent emission usually increases during a burst.
- ii) Photospheric radius expansion adds a confounding spectral effect, likely caused by the changing structure of the photosphere. When this is corrected for, by considering only non-PRE bursts, the spectral fits are much improved.
- iii) Any remaining discrepancy from a blackbody plus variable persistent normalization model is likely to be due to a slight non-blackbody character to the burst emission, rather than spectral shape changes in the persistent emission. This is because changes in the shape of the persistent emission are not detectable in the presence of an additional burst component with the quality of the data currently available. Investigations into the response of the accretion spectrum to a burst transient are desirable.

- iv) I have found no evidence that the persistent emission spectrum changes shape on timescales of up to an hour before or after a burst, and no major changes in its normalization.
- v) For sources that have never shown a radius expansion burst, I obtain a lower bound on their Eddington luminosity from the observation that the persistent emission enhancement for non-PRE bursts seldom exceeds a certain fraction (about 37%) of the Eddington flux. For three sources, this new lower bound was found to be a better constraint than the peak fluxes of the brightest (non-PRE) bursts for those sources.

Because varying the normalization of the persistent flux gives statistically superior spectral fits to fits in which it is held constant, it is a superior way to analyse type I X-ray bursts.

7.2 Difficulties

7.2.1 Degeneracy of spectral components

A potential difficulty is that, if the persistent emission can reach $\sim 80\%$ of the Eddington flux, then this accounts for most of the total flux, and there is little left for the burst component. It seems likely that much of the persistent emission may be re-processed and reflected burst emission. That this should happen is only natural because accretion luminosity is produced primarily at and near the stellar surface and is probably blackbody-like in character. The emitted spectrum is then modified in various ways to produce the persistent emission spectrum I observe. Any blackbody-like emission at the stellar surface will presumably be modified in similar ways; this includes nuclear burning.

It may therefore be that what I identify as persistent emission began as burst emission from thermonuclear burning, but has been modified so that it resembles the persistent emission in spectral shape. However, some non-negligible accretion rate change is still to be expected. An observational challenge is to determine how much apparent persistent emission enhancement is to be identified as modified burst flux.

As mentioned in Chapter 1, half of the gravitational potential energy of the accreting material is released before it reaches the neutron star surface. This radiation itself may irradiate outer portions of the disc. Since the material in the disc, even near the neutron star, moves more quickly than the neutron star surface, this may somewhat reduce the

effectiveness of radiation drag. It follows that, though both halves of the accretion luminosity cause outwards radiation pressure, the transverse drag may be much less effective than the total accretion luminosity would suggest. If the neutron star itself is spinning rapidly, transverse drag will be even less effective.

7.2.2 Magnitude of accretion enhancement

The computational study showed that a significant increase in f_a , up to a factor of 8, is possible, provided that the angular momentum loss caused by radiation drag is efficiently communicated to the interior portions of the disc. I have not been able to reproduce the 37% Eddington accretion rates indicated by the observations, though enhancements up to 20% or so are plausible. Although it is likely that accretion rate can change during a burst, it is unlikely that equating f_a with \dot{M} is a full explanation. As discussed above, it may be that some of the burst emission is modified to look like persistent emission via reflection off the disc (e.g., [Keek et al. 2014](#)) or other processes. Even so, the fact that f_a does not scale linearly with burst flux suggests that a simple instantaneous modification of a seed burst spectrum cannot explain the temporal behaviour of the changing persistent flux.

7.3 Future prospects

7.3.1 Theory

In Chapter 6 I found that decreasing the number of SPH particles from 2.5 million to 250,000 did not significantly change the qualitative behaviour of f_a or angular momentum loss, even though the disc scale height is no longer properly resolved near the star. This fortuitous property makes it possible to explore the parameter space more widely, because the lower-resolution simulations finish much sooner (about four days, compared to a few months for the full resolution runs). I intend to study the effect of increasing the geometric thickness of the disc, of changing the mass of the star, and the effect of selecting a burst profile with a faster rise time. This work, as well as the results of Chapter 6, will be the subject of an upcoming paper.

In the longer term, proper treatment of radiation transport will be desirable, including the thermodynamic response of the disc, rather than the approximations made in this work. The effects of magnetic fields and general relativity also need to be taken into account.

7.3.2 Observational work

Further observational work will have to await the next generation of large X-ray telescopes, such as the *Large Observatory for X-ray Timing* (Feroci et al. 2012) or ATHENA (Nandra et al. 2013). Both of these proposed instruments have similar timing properties to *RXTE* and larger collecting areas by an order of magnitude. Small spectral effects undetectable by *RXTE* will be visible to these instruments. However, these X-ray telescopes are expected to launch in the mid-2020s, at the earliest.

In the meantime, much useful science can be done with current observatories. XMM-Newton is highly suitable for detailed spectral studies, but without high timing resolution, and has been used successfully to probe the accretion environments of X-ray bursters (e.g., Boirin & Parmar 2003). It is particularly suited to the detection of spectral features around 6-7 keV (e.g., Church et al. 2005). This is the location of a fluorescent iron line (e.g., Pandel et al. 2008) arising from reflection from the disc and can potentially break the degeneracy between reflection and enhanced persistent emission.

The *Chandra* telescope can perform spectroscopy at lower energies than *RXTE* (Brinkman et al. 2000). It can study the cooling of the neutron star crust (~ 0.1 keV). It thus provides valuable information about burst ignition conditions, because heating from below is a major influence on the onset of nuclear burning. Absorption by neutral gas (i.e., N_H column density) affects most strongly the energies accessible to *Chandra*, making that telescope ideal for measuring the amount and variability of gas in and around the LMXB system. *Chandra* has also been used to study type I X-ray bursts themselves, in conjunction with *RXTE* with considerable success (e.g., in't Zand et al. 2013).

Multi-instrument observing campaigns of individual bursters can provide coverage that minimises data gaps, as well as covering a larger energy range than the individual instruments. A recent study by Chenevez et al. (2011) of the burster IGR J17473–2721 used four instruments: *RXTE*, *Swift*, *INTEGRAL*, and *AGILE*. The persistent emission of the

source was carefully monitored, and all four instruments detected bursts. The fact that *AGILE* (Tavani et al. 2008), sensitive to X-ray photons between 18 and 60 keV, could successfully identify a burst makes it potentially useful in studies of the corona during a burst, where most of the temporal changes happen at around 30 keV. The recently launched *NuSTAR* (Harrison et al. 2013) is also sensitive to higher energies and is likely to be useful for obtaining data of the “hard tail” of both accretion and burst components of spectra.

The *INTEGRAL* observatory (Winkler et al. 2003) lacks the timing and spectral resolution to effectively study individual bursts in detail, but its larger field of view has allowed it to detect more events than any other telescope, and this provides important data about recurrence times of bursts and therefore something about ignition conditions. *INTEGRAL*'s ability to monitor a large area of sky is backed up by the *MAXI* instrument aboard the International Space Station, which scans almost the whole sky for transient X-ray sources every orbit (Matsuoka et al. 2009), and has already discovered two previously unknown transient sources later identified as X-ray bursters¹. It has also been used in a multi-instrument observing campaign of a non-bursting LMXB (Homan et al. 2014), demonstrating its value in studies of LMXBs in quiescence and in outburst. At the very least, such instruments will ensure that the next generation of X-ray telescopes will not be short of interesting LMXB bursters to study.

¹see <http://www.sron.nl/~jeanz/bursterlist.html>

Appendix A

The table in this Appendix lists N_H values for all sources from which *RXTE* detected radius expansion bursts. These values are drawn from the scientific literature and references are given. I also list Eddington fluxes for these sources, calculated according to the method described in §2.3.

I also list neutron star spin frequencies, where known. These have been calculated either from burst oscillations (BO; e.g., [Watts 2012](#)), or from pulsar timing. I use them in Chapter 3 to test whether the persistent flux enhancements are correlated with the neutron star's rotation period.

Table A.1. N_H values chosen for spectral fits to burst data, Eddington fluxes, and spin periods

Source	N_H (10^{22} cm $^{-2}$)	No. of PRE bursts <i>RXTE</i> (<i>WFC</i>) ^d	Ref. (N_H)	Eddington flux 10^{-9} erg cm $^{-2}$ s $^{-1}$	Spin (Hz)	Method ^c	Ref. (Spin)
4U 0513–40	0.03	4 (3)	1	14.5 ± 3.5			
EXO 0748–676	0.80	5 (0)	2	46.5 ± 4.6	552	BO	27, 28
2S 0918–549	0.35	2 (3)	3	119.2 ± 12.4			
4U 1608–522	0.89	17 (5)	4	172.2 ± 21.8	620	BO	27, 29, 30
4U 1636–536 (high) ^a	0.25	76 (2)	5	73.9 ± 6.8	581	BO	27, 31, 32
4U 1636–536 (low)	0.25	2 ^b (0)	5	41.5 ± 1.2	581	BO	27, 31, 32
MXB 1658–298	0.20	12 (0)	6	17.0 ± 4.0	567	BO	27, 33
XTE 1701–462	2.00	2 (0)	7	43.4 ± 1.4			
4U 1702–429	1.87	5 (3)	8	87.7 ± 4.5	329	BO	27, 34
4U 1705–44	1.90	4 (0)	9	41.0 ± 3.8			
XTE J1710–281	0.40	3 (0)	10	7.1 ± 1.5			
4U 1722–30	0.78	3 (23)	1	61.7 ± 12.4			
4U 1728–34	2.60	94 (0)	11	95.0 ± 8.4	363	BO	27, 35
KS 1731–260	1.30	4 (3)	12	48.6 ± 5.6	524	BO	27, 36, 37
4U 1735–444	0.14	10 (0)	13	34.2 ± 5.6			
GRS 1741.9–2853	11.30	6 (1)	14	35.3 ± 10.9	589	BO	27, 38
1A 1742–294	1.16	2 (0)	8	37.8 ± 1.4			
SLX 1744–300	4.50	4 (3)	15	13.9 ± 3.1			
GX 3+1	1.59	2 (0)	16	60.0 ± 1.4			
SAX J1748.9–2021	0.47	11 (0)	1	38.0 ± 6.0	442	Pulsar	27, 39
EXO 1745–248	3.80	2 (0)	1	69.0 ± 2.8			
4U 1746–37	0.26	3 (0)	1	5.4 ± 0.7			
SAX J1747.0–2853	8.80	11 (2)	17	52.5 ± 7.1			
IGR J17473–2721	3.80	3 (0)	18	113.5 ± 12.1			
IGR J17498–2921	1.20	1 (0)	8	51.6 ± 1.6	401	Pulsar	27, 40
XTE J1759–220	2.84	3 (0)	8	15.7 ± 0.8			
SAX J1750.8–2900	0.90	2 (1)	8	54.1 ± 2.1	601	BO	27, 30, 41
GRS 1747–312	1.39	3 (0)	1	13.4 ± 4.4			
SAX J1808.4–3658	0.12	8 (2)	19	230.1 ± 13.2	401	Pulsar	27, 42
XTE J1810–189	4.20	1 (0)	20	54.2 ± 1.8			
SAX J1810.8–2609	0.35	1 (1)	21	111.2 ± 3.0			
GX 17+2	1.90	2 (0)	22	15.5 ± 0.5			
4U 1820–303	0.16	16 (34)	1	60.5 ± 4.0			
XB 1832–330	0.05	1 (0)	1	33.7 ± 4.4			
HETE J1900.1–2455	0.16	7 (0)	23	123.9 ± 10.6	378	Pulsar	27, 43
Aql X–1	0.40	14 (0)	24	99.6 ± 21.3	550	Pulsar	27, 44

Table A.1 (cont'd)

Source	N_H (10^{22} cm^{-2})	No. of PRE bursts <i>RXTE</i> (<i>WFC</i>) ^d	Ref. (N_H)	Eddington flux $10^{-9} \text{ erg cm}^{-2} \text{ s}^{-1}$	Spin (Hz)	Method ^c	Ref. (Spin)
XB 1916–053	0.32	12 (0)	25	30.6 ± 3.6	270	BO	27, 45
4U 2129+12	0.03	1 (1)	1	40.8 ± 1.6			
Cyg X–2	0.05	8 (0)	26	13.1 ± 2.1			
Ser X–1	0.38	6 (0)	26	29.4 ± 7.1			

^aPRE burst fluxes from 4U 1636–536 are bimodal, separated by a factor of ≈ 1.7 , thought to be due to bursts occurring in different atmospheric compositions (Ebisuzaki & Nakamura 1988, Galloway et al. 2006). Mean Eddington fluxes for both regimes are given.

^bThis includes a tentatively identified PRE burst that is not included in the analysis in the rest of the thesis.

^cBO = Burst Oscillation, Pulsar = Pulsar Timing

^dWFC = *BeppoSAX* Wide Field Camera

References. — 1. Kuulkers et al. (2003); 2. Homan et al. (2003); 3. Juett et al. (2001); 4. Keek et al. (2008); 5. Asai et al. (2000); 6. Oosterbroek, Parmar, Sidoli, in't Zand & Heise (2001); 7. Lin et al. (2009); 8. J. in't Zand; private communication (2012); 9. Piraino et al. (2007); 10. Younes et al. (2009); 11. D'Ai et al. (2006); 12. Cackett et al. (2006); 13. Augusteijn et al. (1998); 14. Sakano et al. (2002); 15. Mori et al. (2005); 16. Oosterbroek, Barret, Guainazzi & Ford (2001); 17. Werner et al. (2004); 18. Altamirano et al. (2008); 19. Wang et al. (2001); 20. Krimm et al. (2008); 21. Natalucci et al. (2000); 22. Farinelli et al. (2007); 23. Campana (2005); 24. Campana & Stella (2003); 25. Church et al. (1998); 26. Asai et al. (2000); 27. Watts (2012); 28. Galloway et al. (2010); 29. Hartman et al. (2003); 30. G08; 31. Strohmayer et al. (1998); 32. Strohmayer & Markwardt (2002); 33. Wijnands et al. (2001); 34. Markwardt et al. (1999); 35. Strohmayer et al. (1996); 36. Smith et al. (1997); 37. Muno et al. (2000); 38. Strohmayer et al. (1997); 39. Altamirano et al. (2008); 40. Papitto et al. (2011); 41. Kaaret et al. (2002); 42. Chakrabarty et al. (2003); 43. Watts et al. (2009); 44. Zhang et al. (1998); 45. Galloway et al. (2001)

Appendix B

To perform the analysis outlined in §5.3 I have repeated the Eddington flux calculations for each source, with the addition of several sources that have radius expansion bursts recorded by other X-ray satellites but none in *RXTE* data. Most of these new sources are neutron stars where *BeppoSAX* recorded PRE bursts but *RXTE* did not, but I have also included a small number of bursts detected by earlier satellites, as listed in Table B.2. Figure B.1 shows the calculation of the Eddington flux for a source that has bursts in all three catalogs.

Table B.1. N_H Values and Eddington Fluxes

Source	Non-PRE Bursts	PRE bursts	SWFC ^a	SBC ^b	N_H 10^{22} cm^{-2}	Reference	F_{Edd} $10^{-9} \text{ erg s}^{-1} \text{ cm}^{-2}$
4U 1746-37	22	3	0	2	0.3	1	5.8 ± 1.4
MXB 1659-298	11	12	0	0	0.2	2	17.0 ± 4.0
Aql X-1	52	12	0	0	0.4	3	99.6 ± 21.3
SAX J1810.8-2609	4	1	1	0	0.3	4	111.2 ± 2.5
IGR J17480-2446	277	0	0	0	0.5	1	...
Cyg X-2	34	7	0	0	0.1	5	13.1 ± 2.1
GS 1826-24	67	0	0	0	0.4	6	...
XTE J1709-267	2	0	0	0	0.4	7	...
4U 1820-30	0	16	34	7	0.2	1	57.1 ± 7.7
SAX 1808.4-3658	1	8	2	0	0.1	8	230.1 ± 13.2
GX 17+2	5	2	0	4	1.9	9	13.3 ± 2.5
Rapid Burster	19	0	0	0	1.7	10	...
IGR J17473-2721	2	0	0	0	3.8	11	113.5 ± 12.1
Cir X-1	13	0	0	0	0.7	12	...
HETE J1900.1-2455	0	7	0	0	0.2	13	123.9 ± 10.6
4U 2129+12	5	1	1	1	0.0	1	40.8 ± 1.6
1M 0836-425	17	0	0	0	2.2	14	...
4U 1735-444	7	10	0	0	0.1	15	34.2 ± 5.6
1A 1744-361	3	0	0	0	4.5	16	...
4U 1722-30	0	3	23	1	0.8	1	61.7 ± 12.4
GX 3+1	1	2	0	1	1.6	17	59.9 ± 0.9
EXO 1745-248	15	2	0	0	3.8	1	69.0 ± 0.2
XTE J1739-285	6	0	0	0	2.0	18	...
2S 0918-549	1	2	3	0	0.3	19	119.2 ± 12.4
Ser X-1	11	6	0	0	0.4	5	29.4 ± 7.1
GRS 1741.9-2853	1	4	1	0	11.3	20	35.3 ± 10.9
2E 1742.9-2929	65	1	0	0	1.2	18	37.8 ± 1.4
4U 1728-34	30	64	0	0	2.6	21	95.0 ± 8.4
XB 1832-330	0	1	0	0	0.1	1	33.7 ± 4.4
IGR J17511-3057	7	0	0	0	0.6	22	...
4U 1608-522	28	18	5	0	0.9	23	167.2 ± 26.0
GRS 1747-312	1	1	0	0	1.4	1	13.4 ± 4.4
XTE 1814-338	27	0	0	0	0.2	18	...
XTE J1723-376	2	0	0	0	7.9	18	...
4U 1254-69	3	0	0	0	0.3	24	...
4U 1702-429	43	5	3	0	1.9	18	87.7 ± 4.5

Table B.1 (cont'd)

Source	Non-PRE Bursts	PRE bursts	SWFC ^a	SBC ^b	N_H 10^{22} cm^{-2}	Reference	F_{Edd} $10^{-9} \text{ erg s}^{-1} \text{ cm}^{-2}$
SAX J1750.8-2900	3	2	1	0	0.9	18	54.1 ± 2.1
XTE J1701-462	1	2	0	0	2.0	25	43.4 ± 1.4
XTE J1810-189	3	1	0	0	4.2	26	54.2 ± 1.8
EXO 0748-676	143	5	0	0	0.8	27	46.5 ± 4.6
4U 1916-053	1	12	0	0	0.3	28	30.6 ± 3.6
XTE 2123-058	3	0	0	0	0.1	29	...
XTE 1759-220	5	3	0	0	2.8	18	15.7 ± 0.8
4U 1636-536	269	76	2	0	0.2	5	72.6 ± 9.1
IGR 17191-2821	5	0	0	0	0.3	30	...
SLX 1744-300	10	0	3	0	4.5	31	13.9 ± 3.1
SAX J1748.9-2021	16	11	0	0	0.5	1	38.0 ± 6.0
4U 1323-62	35	0	0	0	2.4	32	...
KS 1731-260	21	4	3	0	1.3	33	48.6 ± 5.6
XTE J1710-281	37	3	0	0	0.4	34	7.1 ± 1.5
IGR J17498-2921	0	1	0	0	1.2	18	51.6 ± 1.6
SAX J1747.0-2853	6	10	2	0	8.8	35	52.5 ± 7.1
SLX 1735-269	1	0	0	3	0.1	36	49.4 ± 3.9
4U 1705-44	77	4	0	0	1.9	37	41.0 ± 3.8
4U 0513-40	11	4	3	0	0.0	1	14.5 ± 3.5
SAX J1806.5-2215	4	0	0	0	1.0	18	...

References. — 1. Kuulkers et al. (2003); 2. Oosterbroek, Parmar, Sidoli, in't Zand & Heise (2001); 3. Campana & Stella (2003); 4. Natalucci et al. (2000); 5. Asai et al. (2000); 6. in 't Zand et al. (1999); 7. Jonker et al. (2003); 8. Wang et al. (2001); 9. Farinelli et al. (2007); 10. Frogel et al. (1995); 11. Altamirano et al. (2008); 12. Iaria et al. (2005); 13. Campana (2005); 14. Belloni et al. (1993); 15. Augusteijn et al. (1998); 16. Gavriil et al. (2012); 17. Oosterbroek, Barret, Guainazzi & Ford (2001); 18. J. in 't Zand 2014, private communication; 19. Juett et al. (2001); 20. Sakano et al. (2002); 21. D'Aí et al. (2006); 22. Papitto et al. (2010); 23. Keek et al. (2008); 24. Boirin & Parmar (2003); 25. Lin et al. (2009); 26. Krimm et al. (2008); 27. Homan et al. (2003); 28. Church et al. (1998); 29. Hynes et al. (2001); 30. Klein-Wolt et al. (2007); 31. Mori et al. (2005); 32. Church et al. (2005); 33. Cackett et al. (2006); 34. Younes et al. (2009); 35. Werner et al. (2004); 36. David et al. (1997); 37. Piraino et al. (2007)

^aBeppoSAX Wide-Field Camera

^bSupplemental Burst Catalog (see table B.2)

^cNo PRE burst has been observed from this source

Table B.2. Supplemental PRE bursts

Source	Burst Time (MJD)	Instrument	Peak flux ($10^{-9} \text{ erg s}^{-1} \text{ cm}^{-2}$)	Reference
GX 17+2	44749.0	Hakucho	9.3 ± 1.0	1
GX 17+2	44749.5	Hakucho	12.2 ± 1.1	1
GX 17+2	45194.8	Hakucho	10.1 ± 1.1	1
GX 17+2	45197.6	Hakucho	12.8 ± 1.0	1
4U 2129+12	47454.7	Ginga	40.0	2
SLX 1735-269	52897.7	JEM-X	57.6 ± 9.6	3
SLX 1735-269	53108.0	IBIS/INTEGRAL	64.0	4
SLX 1735-269	56267.2	MAXI/GSC	48.0 ± 3.0	5
GX 3+1	53248.8	JEM-X	53.4	6
4U 1722-30	42666.0	OSO-8	60.0	7
4U 1746-37	46317.0	EXOSAT	10.0 ± 1.0	8
4U 1746-37	46318.0	EXOSAT	9.0 ± 1.0	8
4U 1820-303	42550.0	SAS-3/RMC	39.0 ± 2.0	9
4U 1820-303	42940.0	SAS-3/RMC	53.0 ± 4.2	9
4U 1820-303	42941.0	SAS-3/RMC	41.0 ± 4.5	9
4U 1820-303	42942.0	SAS-3/RMC	46.0 ± 3.7	9
4U 1820-303	42942.0	SAS-3/RMC	43.0 ± 3.9	9
4U 1820-303	42944.0	SAS-3/RMC	42.0 ± 2.9	9
4U 1820-303	46296.0	EXOSAT	52.8 ± 1.9	10

References. — 1. Tawara et al. (1984); 2. van Paradijs et al. (1990); 3. Molkov et al. (2005); 4. Negoro et al. (2012); 5. Sguera et al. (2007); 6. Chenevez et al. (2006); 7. Swank et al. (1977); 8. Sztajno et al. (1987); 9. Vacca et al. (1986); 10. Damen et al. (1990)

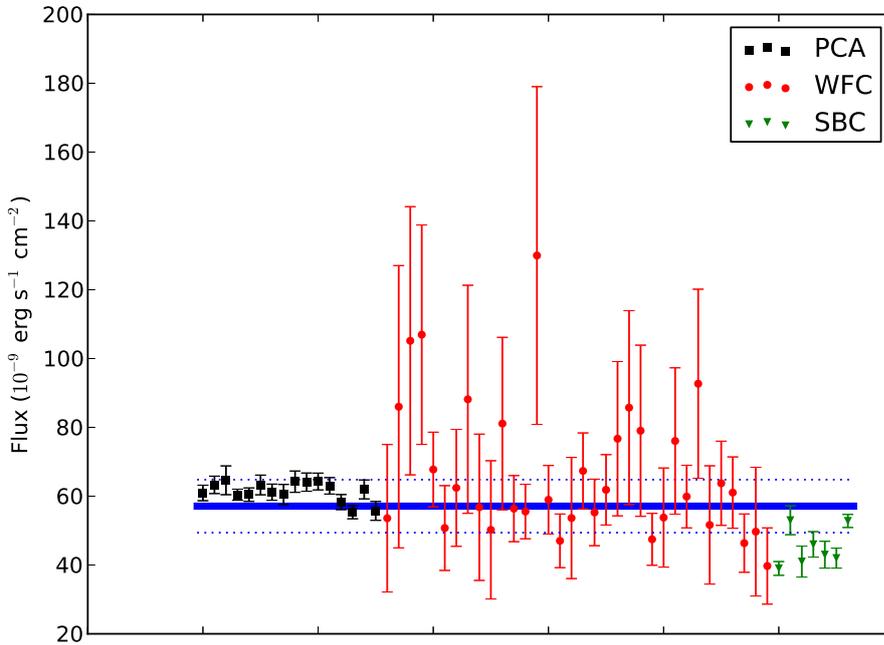


Figure B.1 Measurement of the Eddington flux of 4U 1820–30 from the peak fluxes of its individual radius expansion bursts. The black, red, and green points are measurements from the PCA instruments on *RXTE*, the wide field camera on BeppoSAX, and the Supplementary Burst Catalogue (Table B.2) respectively. The thick blue line is the calculated Eddington flux and the dotted blue lines are the $\pm 1\sigma$ uncertainties on that measurement. Systematic differences between the three catalogs are evident, with *RXTE* bursts consistently higher than those in the supplemental catalog.

Bibliography

- Abramowicz, M. A., Ellis, G. F. R. & Lanza, A. (1990), ‘Relativistic effects in superluminal jets and neutron star winds’, *Astrophysical Journal* **361**, 470–482.
- Adams, D. J. (1980), *Cosmic X-ray astronomy*, Taylor & Francis.
- Alme, M. L. & Wilson, J. R. (1973), ‘X-Ray Emission from a Neutron Star Accreting Material’, *ApJ* **186**, 1015–1026.
- Alpar, M. A., Anderson, P. W., Pines, D. & Shaham, J. (1981), ‘Giant glitches and pinned vorticity in the VELA and other pulsars’, *ApJ* **249**, L29–L33.
- Altamirano, D., Galloway, D. K., Chenevez, J., in’t Zand, J., Kuulkers, E., Degenaar, N., Del Monte, E., Feroci, M., Costa, E., Evangelista, Y., Falanga, M., Markwardt, C., Wijnands, R. & Van Der Klis, M. (2008), ‘Kilohertz QPOs, spectral state transitions and the distance to the neutron star X-ray transient IGR J17473-2721.’, *The Astronomer’s Telegram* **1651**, 1.
- Arnaud, K. A. (1996), XSPEC: The First Ten Years, in G. H. Jacoby & J. Barnes, ed., ‘Astronomical Data Analysis Software and Systems V’, Vol. 101 of *Astronomical Society of the Pacific Conference Series*, p. 17.
- Artymowicz, P. & Lubow, S. H. (1994), ‘Dynamics of binary-disk interaction. 1: Resonances and disk gap sizes’, *ApJ* **421**, 651–667.
- Asai, K., Dotani, T., Nagase, F. & Mitsuda, K. (2000), ‘Iron K Emission Lines in the Energy Spectra of Low-Mass X-Ray Binaries Observed with ASCA’, *ApJS* **131**, 571–591.
- Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., Greenfield, P., Droettboom, M., Bray, E., Aldcroft, T., Davis, M., Ginsburg, A., Price-Whelan, A. M., Kerzendorf, W. E., Conley, A., Crighton, N., Barbary, K., Muna, D., Ferguson, H., Grollier, F., Parikh, M. M., Nair, P. H., Unther, H. M., Deil, C., Woillez, J., Conseil, S., Kramer, R., Turner, J. E. H., Singer, L., Fox, R., Weaver, B. A., Zabalza, V., Edwards, Z. I., Azalee Bostroem, K., Burke, D. J., Casey, A. R., Crawford, S. M., Dencheva, N., Ely, J., Jenness, T., Labrie, K., Lim, P. L., Pierfederici, F., Pontzen, A., Ptak, A., Refsdal, B., Servillat, M. & Streicher, O. (2013), ‘Astropy: A community Python package for astronomy’, *A&A* **558**, A33.
- Augusteijn, T., van der Hooft, F., de Jong, J. A., van Kerkwijk, M. H. & van Paradijs, J. (1998), ‘Phase-resolved spectroscopy of the low-mass X-ray binaries 1636-536/V 801 Arae and 1735-444/V 926 Scorpii’, *A&A* **332**, 561–568.
- Asyali, S. & Joss, P. C. (1982), ‘Thermonuclear processes on accreting neutron stars - A systematic study’, *ApJ* **256**, 637–665.

- Bagnoli, T., in't Zand, J. J. M., Galloway, D. K. & Watts, A. L. (2013), 'Indications for a slow rotator in the Rapid Burster from its thermonuclear bursting behaviour', *MNRAS* **431**, 1947–1955.
- Bagnoli, T., in't Zand, J. J. M., Patruno, A. & Watts, A. L. (2014), 'Double-peaked thermonuclear bursts at the soft-hard state transition in the Rapid Burster', *MNRAS* **437**, 2790–2801.
- Balbus, S. A. & Hawley, J. F. (1998), 'Instability, turbulence, and enhanced transport in accretion disks', *Reviews of Modern Physics* **70**, 1–53.
- Ballantyne, D. R. & Everett, J. E. (2005), 'On the Dynamics of Suddenly Heated Accretion Disks around Neutron Stars', *ApJ* **626**, 364–372.
- Bath, G. T. & Pringle, J. E. (1982), 'The evolution of viscous discs. II - Viscous variations', *MNRAS* **199**, 267–280.
- Bauböck, M., Özel, F., Psaltis, D. & Morsink, S. M. (2015), 'Rotational Corrections to Neutron-star Radius Measurements from Thermal Spectra', *ApJ* **799**, 22.
- Belloni, T., Hasinger, G., Pietsch, W., Mereghetti, S., Bignami, G. F. & Caraveo, P. (1993), 'ROSAT and Optical Observations of Two X-Ray Transients - MX:0836-42 and GS:0834-430', *A&A* **271**, 487.
- Bildsten, L. (1998), Thermonuclear Burning on Rapidly Accreting Neutron Stars, in R. Buccheri, J. van Paradijs, & A. Alpar, ed., 'NATO ASIC Proc. 515: The Many Faces of Neutron Stars.', p. 419.
- Blackburn, J. K. (1995), FTOOLS: A FITS Data Processing and Analysis Software Package, in R. A. Shaw, H. E. Payne & J. J. E. Hayes, eds, 'Astronomical Data Analysis Software and Systems IV', Vol. 77 of *Astronomical Society of the Pacific Conference Series*, p. 367.
- Blissett, R. J. & Cruise, A. M. (1979), 'The restoration of astronomical X-ray spectra', *MNRAS* **186**, 45–57.
- Boirin, L. & Parmar, A. N. (2003), 'Discovery of narrow X-ray absorption features from the low-mass X-ray binary X 1254-690 with XMM-Newton', *A&A* **407**, 1079–1084.
- Boutloukos, S., Miller, M. C. & Lamb, F. K. (2010), 'Super-Eddington Fluxes During Thermonuclear X-ray Bursts', *ApJ* **720**, L15–L19.
- Brinkman, B. C., Gunsing, T., Kaastra, J. S., van der Meer, R., Mewe, R., Paerels, F. B., Raassen, T., van Rooijen, J., Braeuninger, H. W., Burwitz, V., Hartner, G. D., Kettenring, G., Predehl, P., Drake, J. J., Johnson, C. O., Kenter, A. T., Kraft, R. P., Murray, S. S., Ratzlaff, P. W. & Wargelin, B. J. (2000), Description and performance of the low-energy transmission grating spectrometer on board Chandra, in J. E. Truemper & B. Aschenbach, eds, 'X-Ray Optics, Instruments, and Missions III', Vol. 4012 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pp. 81–90.
- Buchler, J. R. & Yueh, W. R. (1976), 'Compton scattering opacities in a partially degenerate electron plasma at high temperatures', *Astrophysical Journal* **210**, 440–446.
- Burger, H. L. & Katz, J. I. (1983), 'The Eddington limit and supercritical accretion. II - Time-dependent calculations', *ApJ* **265**, 393–401.
- Cackett, E. M., Miller, J. M., Ballantyne, D. R., Barret, D., Bhattacharyya, S., Boutelier, M., Miller, M. C., Strohmayer, T. E. & Wijnands, R. (2010), 'Relativistic Lines and Reflection from the Inner Accretion Disks Around Neutron Stars', *ApJ* **720**, 205–225.

- Cackett, E. M., Wijnands, R., Linares, M., Miller, J. M., Homan, J. & Lewin, W. H. G. (2006), ‘Cooling of the quasi-persistent neutron star X-ray transients KS 1731-260 and MXB 1659-29’, *MNRAS* **372**, 479–488.
- Campana, S. (2005), ‘Swift XRT observation of HETE J1900.1-2455’, *The Astronomer’s Telegram* **535**, 1.
- Campana, S. & Stella, L. (2003), ‘The Evolution of the High-Energy Tail in the Quiescent Spectrum of the Soft X-Ray Transient Aquila X-1’, *ApJ* **597**, 474–478.
- Casares, J., Cornelisse, R., Steeghs, D., Charles, P. A., Hynes, R. I., O’Brien, K. & Strohmayer, T. E. (2006), ‘Detection of the irradiated donor in the LMXBs 4U 1636-536 (=V801 Ara) and 4U 1735-444 (=V926 Sco)’, *MNRAS* **373**, 1235–1244.
- Cavecchi, Y., Patruno, A., Haskell, B., Watts, A. L., Levin, Y., Linares, M., Altamirano, D., Wijnands, R. & van der Klis, M. (2011), ‘Implications of Burst Oscillations from the Slowly Rotating Accreting Pulsar IGR J17480-2446 in the Globular Cluster Terzan 5’, *ApJ* **740**, L8.
- Cavecchi, Y., Watts, A. L., Braithwaite, J. & Levin, Y. (2013), ‘Flame propagation on the surfaces of rapidly rotating neutron stars during Type I X-ray bursts’, *MNRAS* **434**, 3526–3541.
- Chaboyer, B., Sarajedini, A. & Armandroff, T. E. (2000), ‘The Age of the Inner Halo Globular Cluster NGC 6652’, *AJ* **120**, 3102–3110.
- Chakrabarty, D., Morgan, E. H., Munro, M. P., Galloway, D. K., Wijnands, R., van der Klis, M. & Markwardt, C. B. (2003), ‘Nuclear-powered millisecond pulsars and the maximum spin frequency of neutron stars’, *Nature* **424**, 42–44.
- Chakrabarty, D., Tomsick, J., Grefenstette, B., Barret, D., Boggs, S. E., Christensen, F., Craig, W. W., Hailey, C. J., Harrison, F., Psaltis, D., Stern, D., Wik, D. R. & Zhang, W. (2014), NuSTAR/XMM-Newton Detection of a Hard Cut-Off in Cen X-4, in ‘American Astronomical Society Meeting Abstracts’, Vol. 223 of *American Astronomical Society Meeting Abstracts*, p. 438.19.
- Chen, Y.-P., Zhang, S., Zhang, S.-N., Ji, L., Torres, D. F., Kretschmar, P., Li, J. & Wang, J.-M. (2013), ‘The Hard X-Ray Behavior of Aql X-1 during Type-I Bursts’, *ApJ* **777**, L9.
- Chenevez, J., Altamirano, D., Galloway, D. K., in’t Zand, J. J. M., Kuulkers, E., Degenaar, N., Falanga, M., Del Monte, E., Evangelista, Y., Feroci, M. & Costa, E. (2011), ‘Puzzling thermonuclear burst behaviour from the transient low-mass X-ray binary IGR J17473-2721’, *MNRAS* **410**, 179–189.
- Chenevez, J., Falanga, M., Brandt, S., Farinelli, R., Frontera, F., Goldwurm, A., in’t Zand, J. J. M., Kuulkers, E. & Lund, N. (2006), ‘Two-phase X-ray burst from GX 3+1 observed by INTEGRAL’, *A&A* **449**, L5–L8.
- Christian, D. J. & Swank, J. H. (1997), ‘The Survey of Low-Mass X-Ray Binaries with the Einstein Observatory Solid-State Spectrometer and Monitor Proportional Counter’, *ApJS* **109**, 177.
- Churazov, E., Gilfanov, M., Forman, W. & Jones, C. (1996), ‘Mapping the Gas Temperature Distribution in Extended X-Ray Sources and Spectral Analysis in the Case of Low Statistics: Application to ASCA Observations of Clusters of Galaxies’, *ApJ* **471**, 673.
- Church, M. J., Parmar, A. N., Balucinska-Church, M., Oosterbroek, T., dal Fiume, D. & Orlandini, M. (1998), ‘Progressive covering in dipping and Comptonization in the spectrum of XB 1916-053 from the BeppoSAX observation’, *A&A* **338**, 556–562.

- Church, M. J., Reed, D., Dotani, T., Bahucińska-Church, M. & Smale, A. P. (2005), ‘Discovery of absorption features of the accretion disc corona and systematic acceleration of the X-ray burst rate in XB1323-619’, *MNRAS* **359**, 1336–1344.
- Cohn, H. N., Lugger, P. M., Grindlay, J. E. & Edmonds, P. D. (2002), ‘Hubble Space Telescope/NICMOS Observations of Terzan 5: Stellar Content and Structure of the Core’, *ApJ* **571**, 818–829.
- Coroniti, F. V. (1981), ‘On the magnetic viscosity in Keplerian accretion disks’, *ApJ* **244**, 587–599.
- Cumming, A., Zweibel, E. & Bildsten, L. (2001), ‘Magnetic Screening in Accreting Neutron Stars’, *ApJ* **557**, 958–966.
- D’Aí, A., di Salvo, T., Iaria, R., Méndez, M., Burderi, L., Lavagetto, G., Lewin, W. H. G., Robba, N. R., Stella, L. & van der Klis, M. (2006), ‘The iron K-shell features of MXB 1728-34 from a simultaneous Chandra-RXTE observation’, *A&A* **448**, 817–822.
- Damen, E., Magnier, E., Lewin, W. H. G., Tan, J., Penninx, W. & van Paradijs, J. (1990), ‘X-ray bursts with photospheric radius expansion and the gravitational redshift of neutron stars’, *A&A* **237**, 103–109.
- D’Angelo, C. R. & Spruit, H. C. (2012), ‘Accretion discs trapped near corotation’, *MNRAS* **420**, 416–429.
- David, P., Goldwurm, A., Murakami, T., Paul, J., Laurent, P. & Goldoni, P. (1997), ‘ASCA X-ray observations of the Galactic bulge source SLX 1735-269.’, *A&A* **322**, 229–233.
- Day, C. S. R. & Tawara, Y. (1990), ‘The Rising Phase of X-Ray Bursts from MXB:1728-34 - Constraints on the Spreading of Thermonuclear Flashes Over the Surface of the Neutron Star’, *MNRAS* **245**, 31P.
- Dorman, B. & Arnaud, K. A. (2001), Redesign and Reimplementation of XSPEC, in F. R. Harnden Jr., F. A. Primini, & H. E. Payne, ed., ‘Astronomical Data Analysis Software and Systems X’, Vol. 238 of *Astronomical Society of the Pacific Conference Series*, p. 415.
- Ebisuzaki, T., Hanawa, T. & Sugimoto, D. (1983), ‘Mass loss from the neutron stars associated with X-ray bursts’, *PASJ* **35**, 17–32.
- Ebisuzaki, T. & Nakamura, N. (1988), ‘The difference in hydrogen abundance between two classes of type I X-ray bursts’, *ApJ* **328**, 251–255.
- Ebisuzaki, T., Sugimoto, D. & Hanawa, T. (1984), ‘Are X-ray bursts really of super-Eddington luminosities?’, *PASJ* **36**, 551–566.
- Fancy, M. K. (2013), ‘Remarks on the Efficacy of Sunbeams in Bringing Small Objects to Rest’, *IJSC* **9**, 221–247.
- Farinelli, R., Titarchuk, L. & Frontera, F. (2007), ‘The Hard X-Ray Tails in Neutron Star Low-Mass X-Ray Binaries: BeppoSAX Observations and Possible Theoretical Explanation of the Case of GX 17+2’, *ApJ* **662**, 1167–1172.
- Feroci, M., Stella, L., van der Klis, M., Courvoisier, T. J.-L., Hernanz, M., Hudec, R., Santangelo, A., Walton, D., Zdziarski, A. & et. al. (2012), ‘The Large Observatory for X-ray Timing (LOFT)’, *Experimental Astronomy* **34**, 415–444.
- Fisker, J. L., Tan, W., Görres, J., Wiescher, M. & Cooper, R. L. (2007), ‘The $^{15}\text{O}(\alpha,\gamma)^{19}\text{Ne}$ Reaction Rate and the Stability of Thermonuclear Burning on Accreting Neutron Stars’, *ApJ* **665**, 637–641.

- Fox, D. W., Lewin, W. H. G., Rutledge, R. E., Morgan, E. H., Guerriero, R., Bildsten, L., van der Klis, M., van Paradijs, J., Moore, C. B., Dotani, T. & Asai, K. (2001), ‘Search for millisecond periodicities in type I X-ray bursts of the Rapid Burster’, *MNRAS* **321**, 776–782.
- Frank, J., King, A. & Raine, D. J. (2002), *Accretion Power in Astrophysics: Third Edition*, Cambridge University Press.
- Frogel, J. A., Kuchinski, L. E. & Tiede, G. P. (1995), ‘Infrared array photometry of metal rich globular clusters. 2: Liller 1—the most metal rich cluster?’, *AJ* **109**, 1154–1168.
- Fujimoto, M. Y., Hanawa, T. & Miyaji, S. (1981), ‘Shell flashes on accreting neutron stars and X-ray bursts’, *ApJ* **247**, 267–278.
- Galloway, D. K., Chakrabarty, D., Munro, M. P. & Savov, P. (2001), ‘Discovery of a 270 Hertz X-Ray Burst Oscillation in the X-Ray Dipper 4U 1916-053’, *ApJ* **549**, L85–L88.
- Galloway, D. K. & Lampe, N. (2012), ‘On the Consistency of Neutron-star Radius Measurements from Thermonuclear Bursts’, *ApJ* **747**, 75.
- Galloway, D. K., Lin, J., Chakrabarty, D. & Hartman, J. M. (2010), ‘Discovery of a 552 Hz Burst Oscillation in the Low-Mass X-Ray Binary EXO 0748-676’, *ApJ* **711**, L148–L151.
- Galloway, D. K., Munro, M. P., Hartman, J. M., Psaltis, D. & Chakrabarty, D. (2008), ‘Thermonuclear (Type I) X-Ray Bursts Observed by the Rossi X-Ray Timing Explorer’, *ApJS* **179**, 360–422.
- Galloway, D. K., Psaltis, D., Munro, M. P. & Chakrabarty, D. (2006), ‘Eddington-limited X-Ray Bursts as Distance Indicators. II. Possible Compositional Effects in Bursts from 4U 1636-536’, *ApJ* **639**, 1033–1038.
- Galloway, D. K., Yao, Y., Marshall, H., Misanovic, Z. & Weinberg, N. (2010b), ‘Radius-expansion Burst Spectra from 4U 1728-34: An Ultracompact Binary?’, *ApJ* **724**, 417–424.
- Gavril, F. P., Strohmayer, T. E. & Bhattacharyya, S. (2012), ‘An Fe XXIV Absorption Line in the Persistent Spectrum of the Dipping Low-mass X-Ray Binary 1A 1744-361’, *ApJ* **753**, 2.
- Gingold, R. A. & Monaghan, J. J. (1977), ‘Smoothed particle hydrodynamics - Theory and application to non-spherical stars’, *MNRAS* **181**, 375–389.
- Gottwald, M., Haberl, F., Parmar, A. N. & White, N. E. (1986), ‘The bursting behavior of the transient X-ray burst source EXO 0748-676 - A dependence between the X-ray burst properties and the strength of the persistent emission’, *Astrophysical Journal* **308**, 213–224.
- Gregory, P. C. (2005), *Bayesian Logical Data Analysis for the Physical Sciences: A Comparative Approach with ‘Mathematica’ Support*, Cambridge University Press.
- Grindlay, J., Gursky, H., Schnopper, H., Parsignault, D. R., Heise, J., Brinkman, A. C. & Schrijver, J. (1976), ‘Discovery of intense X-ray bursts from the globular cluster NGC 6624’, *ApJ* **205**, L127–L130.
- Guess, A. W. (1962), ‘Poynting-Robertson effect for a spherical source of radiation’, *ApJ* **135**, 855–866.
- Güver, T., Özel, F., Cabrera-Lavers, A. & Wroblewski, P. (2010), ‘The Distance, Mass, and Radius of the Neutron Star in 4U 1608-52’, *ApJ* **712**, 964–973.

- Güver, T., Psaltis, D. & Özel, F. (2012), ‘Systematic Uncertainties in the Spectroscopic Measurements of Neutron-star Masses and Radii from Thermonuclear X-Ray Bursts. I. Apparent Radii’, *ApJ* **747**, 76.
- Haberl, F., Stella, L., White, N. E., Gottwald, M. & Priedhorsky, W. C. (1987), ‘EXOSAT observations of double-peaked bursts with radius expansion from 4U/MXB 1820-30’, *ApJ* **314**, 266–271.
- Hanawa, T. & Sugimoto, D. (1982), ‘Helium Shell Flash on Accreting Neutron Stars - Effects of Hydrogen-Rich Envelope and Recurrence of X-Ray Bursts’, *PASJ* **34**, 1.
- Harris, W. E. (1996), ‘A Catalog of Parameters for Globular Clusters in the Milky Way’, *AJ* **112**, 1487.
- Harrison, F. A., Craig, W. W., Christensen, F. E., Hailey, C. J., Zhang, W. W., Boggs, S. E., Stern, D., Cook, W. R., Forster, K., Giommi, P., Grefenstette, B. W., Kim, Y., Kitaguchi, T., Koglin, J. E., Madsen, K. K., Mao, P. H., Miyasaka, H., Mori, K., Perri, M., Pivovarov, M. J., Puccetti, S., Rana, V. R., Westergaard, N. J., Willis, J., Zoglauer, A., An, H., Bachetti, M., Barrière, N. M., Bellm, E. C., Bhalerao, V., Brejnholt, N. F., Fuerst, F., Liebe, C. C., Markwardt, C. B., Nynka, M., Vogel, J. K., Walton, D. J., Wik, D. R., Alexander, D. M., Cominsky, L. R., Hornschemeier, A. E., Hornstrup, A., Kaspi, V. M., Madejski, G. M., Matt, G., Molendi, S., Smith, D. M., Tomsick, J. A., Ajello, M., Ballantyne, D. R., Baloković, M., Barret, D., Bauer, F. E., Blandford, R. D., Brandt, W. N., Brenneman, L. W., Chiang, J., Chakrabarty, D., Chenevez, J., Comastri, A., Dufour, F., Elvis, M., Fabian, A. C., Farrah, D., Fryer, C. L., Gotthelf, E. V., Grindlay, J. E., Helfand, D. J., Krivonos, R., Meier, D. L., Miller, J. M., Natalucci, L., Ogle, P., Ofek, E. O., Ptak, A., Reynolds, S. P., Rigby, J. R., Tagliaferri, G., Thorsett, S. E., Treister, E. & Urry, C. M. (2013), ‘The Nuclear Spectroscopic Telescope Array (NuSTAR) High-energy X-Ray Mission’, *ApJ* **770**, 103.
- Hartman, J. M., Chakrabarty, D., Galloway, D. K., Munro, M. P., Savov, P., Mendez, M., van Straaten, S. & Di Salvo, T. (2003), Discovery of 619 Hz Thermonuclear Burst Oscillations in the Low-Mass X-Ray Binary 4U 1608-52, in ‘AAS/High Energy Astrophysics Division #7’, Vol. 35 of *Bulletin of the American Astronomical Society*, p. 865.
- Hasinger, G. & van der Klis, M. (1989), ‘Two patterns of correlated X-ray timing and spectral behaviour in low-mass X-ray binaries’, *A&A* **225**, 79–96.
- Hesley, J. N., Janes, K. A., Zinn, R., Demarque, P., Da Costa, G. S. & Christian, C. A. (2000), ‘Hubble Space Telescope Photometry of the Metal-rich Globular Clusters NGC 6624 and NGC 6637’, *AJ* **120**, 879–893.
- Heger, A., Cumming, A., Galloway, D. K. & Woosley, S. E. (2007), ‘Models of Type I X-Ray Bursts from GS 1826-24: A Probe of rp-Process Hydrogen Burning’, *ApJ* **671**, L141–L144.
- Heger, A., Cumming, A. & Woosley, S. E. (2007), ‘Millihertz Quasi-periodic Oscillations from Marginally Stable Nuclear Burning on an Accreting Neutron Star’, *ApJ* **665**, 1311–1320.
- Hernquist, L. & Katz, N. (1989), ‘TREESPH - A unification of SPH with the hierarchical tree method’, *ApJS* **70**, 419–446.
- Homan, J., Fridriksson, J. K., Wijnands, R., Cackett, E. M., Degenaar, N., Linares, M., Lin, D. & Remillard, R. A. (2014), ‘A Strongly Heated Neutron Star in the Transient Z Source MAXI J0556-332’, *ApJ* **795**, 131.
- Homan, J., van der Klis, M., Wijnands, R., Belloni, T., Fender, R., Klein-Wolt, M., Casella, P., Méndez, M., Gallo, E., Lewin, W. H. G. & Gehrels, N. (2007), ‘Rossi X-Ray Timing Explorer Observations of the First Transient Z Source XTE J1701-462: Shedding New Light on Mass Accretion in Luminous Neutron Star X-Ray Binaries’, *ApJ* **656**, 420–430.

- Homan, J., Wijnands, R. & van den Berg, M. (2003), ‘XMM-Newton light curves of the low-mass X-ray binary EXO 0748-676: Dips, eclipses, and bursts’, *A&A* **412**, 799–812.
- Hynes, R. I., Charles, P. A., Haswell, C. A., Casares, J., Zurita, C. & Serra-Ricart, M. (2001), ‘Optical studies of the X-ray transient XTE J2123-058 - II. Phase-resolved spectroscopy’, *MNRAS* **324**, 180–190.
- Iaria, R., Spanò, M., Di Salvo, T., Robba, N. R., Burderi, L., Fender, R., van der Klis, M. & Frontera, F. (2005), ‘On the Soft Excess in the X-Ray Spectrum of Circinus X-1: Revisitation of the Distance to Circinus X-1’, *ApJ* **619**, 503–516.
- in ’t Zand, J. J. M., Heise, J., Kuulkers, E., Bazzano, A., Cocchi, M. & Ubertini, P. (1999), ‘Broad-band X-ray measurements of GS 1826-238’, *A&A* **347**, 891–896.
- Inogamov, N. A. & Sunyaev, R. A. (2010), ‘Spread of matter over a neutron-star surface during disk accretion: Deceleration of rapid rotation’, *Astronomy Letters* **36**, 848–894.
- in’t Zand, J. J. M., Cumming, A., Triemstra, T. L., Mateijsen, R. A. D. A. & Bagnoli, T. (2014), ‘The cooling rate of neutron stars after thermonuclear shell flashes’, *A&A* **562**, A16.
- in’t Zand, J. J. M., Galloway, D. K. & Ballantyne, D. R. (2011), ‘Achromatic late-time variability in thermonuclear X-ray bursts. An accretion disk disrupted by a nova-like shell?’, *A&A* **525**, A111+.
- in’t Zand, J. J. M., Galloway, D. K., Marshall, H. L., Ballantyne, D. R., Jonker, P. G., Paerels, F. B. S., Palmer, D. M., Patruno, A. & Weinberg, N. N. (2013), ‘A bright thermonuclear X-ray burst simultaneously observed with Chandra and RXTE’, *A&A* **553**, A83.
- in’t Zand, J. J. M., Keek, L. & Cavecchi, Y. (2014), ‘Relativistic outflow from two thermonuclear shell flashes on neutron stars’, *A&A* **568**, A69.
- in’t Zand, J. J. M. & Weinberg, N. N. (2010), ‘Evidence of heavy-element ashes in thermonuclear X-ray bursts with photospheric superexpansion’, *A&A* **520**, A81+.
- Jahoda, K., Markwardt, C. B., Radeva, Y., Rots, A. H., Stark, M. J., Swank, J. H., Strohmayer, T. E. & Zhang, W. (2006), ‘Calibration of the Rossi X-Ray Timing Explorer Proportional Counter Array’, *ApJS* **163**, 401–423.
- Ji, L., Zhang, S., Chen, Y.-P., Zhang, S.-N., Kretschmar, P., Wang, J.-M. & Li, J. (2014), ‘Possible hard X-ray shortages in bursts from KS 1731-260 and 4U 1705-44’, *A&A* **564**, A20.
- Ji, L., Zhang, S., Chen, Y., Zhang, S.-N., Torres, D. F., Kretschmar, P., Chernyakova, M., Li, J. & Wang, J.-M. (2013), ‘X-ray bursts as a probe of the corona: the case of XRB 4U 1636-536’, *MNRAS* **432**, 2773–2778.
- Ji, L., Zhang, S., Chen, Y., Zhang, S.-N., Torres, D. F., Kretschmar, P. & Li, J. (2014), ‘The Hard X-Ray Shortages Prompted by the Clock Bursts in GS 1826-238’, *ApJ* **782**, 40.
- Jonker, P. G., Méndez, M., Nelemans, G., Wijnands, R. & van der Klis, M. (2003), ‘Chandra observations of the neutron star soft X-ray transient RX J170930.2 - 263927 returning to quiescence’, *MNRAS* **341**, 823–831.
- Jonker, P. G. & Nelemans, G. (2004), ‘The distances to Galactic low-mass X-ray binaries: consequences for black hole luminosities and kicks’, *MNRAS* **354**, 355–366.
- Joss, P. C. (1978), ‘Helium-burning flashes on an accreting neutron star - A model for X-ray burst sources’, *Astrophysical Journal Letters* **225**, L123–L127.

- Joss, P. C. & Melia, F. (1987), ‘Quasi-static winds from neutron stars’, *Astrophysical Journal* **312**, 700–710.
- Juett, A. M., Psaltis, D. & Chakrabarty, D. (2001), ‘Ultracompact X-Ray Binaries with Neon-rich Degenerate Donors’, *ApJ* **560**, L59–L63.
- Kaaret, P., in ’t Zand, J. J. M., Heise, J. & Tomsick, J. A. (2002), ‘Discovery of Millisecond Variability in the Neutron Star X-Ray Transient SAX J1750.8-2900’, *ApJ* **575**, 1018–1024.
- Kato, M. (1983), ‘Neutron Star Wind’, *Publications of the Astronomical Society of Japan* **35**, 33–46.
- Keek, L., Ballantyne, D. R., Kuulkers, E. & Strohmayer, T. E. (2014), ‘Characterizing the Evolving X-Ray Spectral Features during a Superburst from 4U 1636-536’, *ApJ* **789**, 121.
- Keek, L., in’t Zand, J. J. M., Kuulkers, E., Cumming, A., Brown, E. F. & Suzuki, M. (2008), ‘First superburst from a classical low-mass X-ray binary transient’, *A&A* **479**, 177–188.
- Kendall, M. (1948), *Rank correlation methods*, Griffin, London.
- King, A. R., Pringle, J. E. & Livio, M. (2007), ‘Accretion disc viscosity: how big is alpha?’, *MNRAS* **376**, 1740–1746.
- King, A. R. & Watson, M. G. (1986), ‘The shortest period binary star?’, *Nature* **323**, 105.
- Klein-Wolt, M., Maitra, D., Wijnands, R., Swank, J. H., Markwardt, C. B. & Bailyn, C. (2007), ‘An improvement of the position of IGR J17191-2821’, *The Astronomer’s Telegram* **1070**, 1.
- Kluźniak, W. & Wagoner, R. V. (1985), ‘Evolution of the innermost stable orbits around accreting neutron stars’, *ApJ* **297**, 548–554.
- Krimm, H. A., Kennea, J. & Tueller, J. (2008), ‘Swift/XRT reports an improved position for XTE J1810-189’, *The Astronomer’s Telegram* **1432**, 1.
- Kuulkers, E., den Hartog, P. R., in’t Zand, J. J. M., Verbunt, F. W. M., Harris, W. E. & Cocchi, M. (2003), ‘Photospheric radius expansion X-ray bursts as standard candles’, *A&A* **399**, 663–680.
- Kuulkers, E., Homan, J., van der Klis, M., Lewin, W. H. G. & Méndez, M. (2002), ‘X-ray bursts at extreme mass accretion rates from GX 17+2’, *A&A* **382**, 947–973.
- Lamb, F. K. & Miller, M. C. (1995), ‘Critical radiation fluxes and luminosities of black holes and relativistic stars’, *ApJ* **439**, 828–845.
- Lattimer, J. M. & Prakash, M. (2007), ‘Neutron star observations: Prognosis for equation of state constraints’, *Phys. Rep.* **442**, 109–165.
- Levine, A. M., Bradt, H., Cui, W., Jernigan, J. G., Morgan, E. H., Remillard, R., Shirey, R. E. & Smith, D. A. (1996), ‘First Results from the All-Sky Monitor on the Rossi X-Ray Timing Explorer’, *ApJ* **469**, L33.
- Lewin, W. H. G., Hoffman, J. A., Doty, J., Clark, G. W., Swank, J. H., Becker, R. H., Pravdo, S. H. & Serlemitsos, P. J. (1977), ‘Galactic distribution of X-ray burst sources’, *Nature* **267**, 28–30.
- Lewin, W. H. G. & Joss, P. C. (1981), ‘X-ray bursters and the X-ray sources of the galactic bulge’, *Space Science Reviews* **28**, 3–87.
- Lewin, W. H. G., van Paradijs, J. & Taam, R. E. (1993), ‘X-Ray Bursts’, *Space Sci. Rev.* **62**, 223–389.

- Lin, D., Remillard, R. A. & Homan, J. (2009), ‘Spectral States of XTE J1701 - 462: Link Between Z and Atoll Sources’, *ApJ* **696**, 1257–1277.
- Linares, M., Bahramian, A., Heinke, C., Wijnands, R., Patruno, A., Altamirano, D., Homan, J., Bogdanov, S. & Pooley, D. (2014), ‘The neutron star transient and millisecond pulsar in M28: from sub-luminous accretion to rotation-powered quiescence’, *MNRAS* **438**, 251–261.
- Liu, Q. Z., van Paradijs, J. & van den Heuvel, E. P. J. (2001), ‘A catalogue of low-mass X-ray binaries’, *A&A* **368**, 1021–1054.
- Liu, Q. Z., van Paradijs, J. & van den Heuvel, E. P. J. (2007), ‘A catalogue of low-mass X-ray binaries in the Galaxy, LMC, and SMC (Fourth edition)’, *A&A* **469**, 807–810.
- Lodato, G. & Price, D. J. (2010), ‘On the diffusive propagation of warps in thin accretion discs’, *MNRAS* **405**, 1212–1226.
- London, R. A., Howard, W. M. & Taam, R. E. (1984), ‘The spectra of X-ray bursting neutron stars’, *ApJ* **287**, L27–L30.
- London, R. A., Taam, R. E. & Howard, W. M. (1986), ‘Model atmospheres for X-ray bursting neutron stars’, *ApJ* **306**, 170–182.
- Lucy, L. B. (1977), ‘A numerical approach to the testing of the fission hypothesis’, *AJ* **82**, 1013–1024.
- Lyu, M., Méndez, M., Sanna, A., Homan, J., Belloni, T. & Hiemstra, B. (2014), ‘Iron-line and continuum variations in the XMM-Newton and Suzaku spectra of the neutron-star low-mass X-ray binary 4U 1636-53’, *MNRAS* **440**, 1165–1178.
- Maccarone, T. J. & Coppi, P. S. (2003), ‘Spectral fits to the 1999 Aql X-1 outburst data’, *A&A* **399**, 1151–1157.
- Madej, J. (1991), ‘Model atmospheres and X-ray spectra of bursting neutron stars’, *ApJ* **376**, 161–176.
- Madej, J., Joss, P. C. & Różańska, A. (2004), ‘Model Atmospheres and X-Ray Spectra of Bursting Neutron Stars: Hydrogen-Helium Comptonized Spectra’, *ApJ* **602**, 904–912.
- Markwardt, C. B., Strohmayer, T. E. & Swank, J. H. (1999), ‘Observation of Kilohertz Quasi-periodic Oscillations from the Atoll Source 4U 1702-429 by the Rossi X-Ray Timing Explorer’, *ApJ* **512**, L125–L129.
- Masetti, N. (2002), ‘A change in the outburst recurrence time of the Rapid Burster’, *A&A* **381**, L45–L48.
- Matsuoka, M., Kawasaki, K., Ueno, S., Tomida, H., Kohama, M., Suzuki, M., Adachi, Y., Ishikawa, M., Mihara, T., Sugizaki, M., Isobe, N., Nakagawa, Y., Tsunemi, H., Miyata, E., Kawai, N., Kataoka, J., Morii, M., Yoshida, A., Negoro, H., Nakajima, M., Ueda, Y., Chujo, H., Yamaoka, K., Yamazaki, O., Nakahira, S., You, T., Ishiwata, R., Miyoshi, S., Eguchi, S., Hiroi, K., Katayama, H. & Ebisawa, K. (2009), ‘The MAXI Mission on the ISS: Science and Instruments for Monitoring All-Sky X-Ray Images’, *PASJ* **61**, 999–.
- Maxima (2012), ‘Maxima, a Computer Algebra System. Version 5.28.0’.
URL: <http://maxima.sourceforge.net>
- Mihalas, D. (1978), *Stellar atmospheres, 2nd edition*, W. H. Freeman & Company.
- Miller, M. C. & Lamb, F. K. (1993), ‘Effect of radiation forces on disk accretion by weakly magnetic neutron stars’, *ApJ* **413**, L43–L46.

- Miller, M. C. & Lamb, F. K. (1996), ‘Motion of Accreting Matter near Luminous Slowly Rotating Relativistic Stars’, *ApJ* **470**, 1033.
- Molkov, S., Revnivtsev, M., Lutovinov, A. & Sunyaev, R. (2005), ‘INTEGRAL detection of a long powerful burst from SLX 1735-269’, *A&A* **434**, 1069–1075.
- Monaghan, J. J. (1992), ‘Smoothed particle hydrodynamics’, *ARA&A* **30**, 543–574.
- Monaghan, J. J. (2005), ‘Smoothed particle hydrodynamics’, *Reports on Progress in Physics* **68**, 1703–1759.
- Mori, H., Maeda, Y., Pavlov, G. G., Sakano, M. & Tsuboi, Y. (2005), ‘XMM-Newton observations of the Mouse, SLX 1744 299 and SLX 1744 300’, *Advances in Space Research* **35**, 1137–1141.
- Muno, M. P., Chakrabarty, D., Galloway, D. K. & Savov, P. (2001), ‘Millisecond Oscillations and Photospheric Radius Expansion in Thermonuclear X-Ray Bursts’, *ApJ* **553**, L157–L160.
- Muno, M. P., Fox, D. W., Morgan, E. H. & Bildsten, L. (2000), ‘Nearly Coherent Oscillations in Type I X-Ray Bursts from KS 1731-260’, *ApJ* **542**, 1016–1033.
- Murray, J. R. (1996), ‘SPH simulations of tidally unstable accretion discs in cataclysmic variables’, *MNRAS* **279**, 402–414.
- Nandra, K., Barret, D., Barcons, X., Fabian, A., den Herder, J.-W., Piro, L., Watson, M., Adami, C., Aird, J., Afonso, J. M. & et al. (2013), ‘The Hot and Energetic Universe: A White Paper presenting the science theme motivating the Athena+ mission’, *ArXiv e-prints*.
- Natalucci, L., Bazzano, A., Cocchi, M., Ubertini, P., Heise, J., Kuulkers, E., in’t Zand, J. J. M. & Smith, M. J. S. (2000), ‘SAX J1810.8-2609: A New Hard X-Ray Bursting Transient’, *ApJ* **536**, 891–895.
- Negoro, H., Asada, M., Serino, M., Nakahira, S., Morii, M., Ogawa, Y., Ueno, S., Tomida, H., Ishikawa, M., Yamaoka, K., Kimura, M., Mihara, T., Sugizaki, M., Morihana, K., Yamamoto, T., Sugimoto, J., Takagi, T., Matsuoka, M., Kawai, N., Usui, R., Ishikawa, K., Yoshida, A., Sakamoto, T., Nakano, Y., Tsunemi, H., Sasaki, M., Nakajima, M., Ueda, Y., Hiroi, K., Shidatsu, M., Sato, R., Kawamuro, T., Tsuboi, Y., Yamauchi, M. H. M., Nishimura, Y., Hanayama, T. & Yoshidom, K. (2012), ‘MAXI/GSC detection of a possible superburst from SLX 1735-269’, *The Astronomer’s Telegram* **4622**, 1.
- Oosterbroek, T., Barret, D., Guainazzi, M. & Ford, E. C. (2001), ‘Simultaneous BeppoSAX and RXTE observations of the X-ray burst sources GX 3+1 and Ser X-1’, *A&A* **366**, 138–145.
- Oosterbroek, T., Parmar, A. N., Sidoli, L., in’t Zand, J. J. M. & Heise, J. (2001), ‘BeppoSAX observation of the eclipsing dipping X-ray binary X 1658-298’, *A&A* **376**, 532–542.
- Orlandini, M., Frontera, F., Masetti, N., Sguera, V. & Sidoli, L. (2012), ‘BeppoSAX Observations of the X-Ray Pulsar MAXI J1409-619 in Low State: Discovery of Cyclotron Resonance Features’, *ApJ* **748**, 86.
- Ostriker, J. P., Rees, M. J. & Silk, J. (1970), ‘Some Observable Consequences of Accretion by Defunct Pulsars’, *Astrophys. Lett.* **6**, 179.
- Özel, F. (2006), ‘Soft equations of state for neutron-star matter ruled out by EXO 0748 - 676’, *Nature* **441**, 1115–1117.

- Özel, F., Gould, A. & Güver, T. (2012), ‘The Mass and Radius of the Neutron Star in the Bulge Low-mass X-Ray Binary KS 1731-260’, *ApJ* **748**, 5.
- Özel, F., Güver, T. & Psaltis, D. (2009), ‘The Mass and Radius of the Neutron Star in EXO 1745-248’, *ApJ* **693**, 1775–1779.
- Paczynski, B. (1983*a*), ‘A one-zone model for shell flashes on accreting compact stars’, *Astrophysical Journal* **264**, 282–295.
- Paczynski, B. (1983*b*), ‘Models of X-ray bursters with radius expansion’, *Astrophysical Journal* **267**, 315–321.
- Paczynski, B. & Anderson, N. (1986), ‘Models of extended relativistic envelopes of neutron stars’, *ApJ* **302**, 1–10.
- Paizis, A., Farinelli, R., Titarchuk, L., Courvoisier, T. J.-L., Bazzano, A., Beckmann, V., Frontera, F., Goldoni, P., Kuulkers, E., Mereghetti, S., Rodriguez, J. & Vilhu, O. (2006), ‘Average hard X-ray emission from NS LMXBs: observational evidence of different spectral states in NS LMXBs’, *A&A* **459**, 187–197.
- Pandel, D., Kaaret, P. & Corbel, S. (2008), ‘Relativistic Iron Line Emission from the Neutron Star Low-mass X-Ray Binary 4U 1636-536’, *ApJ* **688**, 1288–1294.
- Papitto, A., Bozzo, E., Ferrigno, C., Belloni, T., Burderi, L., di Salvo, T., Riggio, A., D’Aì, A. & Iaria, R. (2011), ‘The discovery of the 401 Hz accreting millisecond pulsar IGR J17498-2921 in a 3.8 h orbit’, *A&A* **535**, L4.
- Papitto, A., Riggio, A., di Salvo, T., Burderi, L., D’Aì, A., Iaria, R., Bozzo, E. & Menna, M. T. (2010), ‘The X-ray spectrum of the newly discovered accreting millisecond pulsar IGR J17511-3057’, *MNRAS* **407**, 2575–2588.
- Peille, P., Olive, J.-F. & Barret, D. (2014), ‘Probing X-ray burst - accretion disk interaction in low mass X-ray binaries through kilohertz quasiperiodic oscillations’, *A&A* **567**, A80.
- Piraino, S., Santangelo, A., di Salvo, T., Kaaret, P., Horns, D., Iaria, R. & Burderi, L. (2007), ‘BeppoSAX observation of 4U 1705-44: detection of hard X-ray emission in the soft state’, *A&A* **471**, L17–L20.
- Podsiadlowski, P., Rappaport, S. & Pfahl, E. D. (2002), ‘Evolutionary Sequences for Low- and Intermediate-Mass X-Ray Binaries’, *ApJ* **565**, 1107–1133.
- Popham, R. & Sunyaev, R. (2001), ‘Accretion Disk Boundary Layers around Neutron Stars: X-Ray Production in Low-Mass X-Ray Binaries’, *ApJ* **547**, 355–383.
- Poynting, J. H. (1904), ‘Radiation in the Solar System: Its Effect on Temperature and Its Pressure on Small Bodies’, *Royal Society of London Philosophical Transactions Series A* **202**, 525–552.
- Press, W. (1996), ‘Understanding data better with Bayesian and global statistical methods’, *Unsolved Problems in Astrophysics* pp. 49–60.
- Price, D. J. (2007), ‘splash: An Interactive Visualisation Tool for Smoothed Particle Hydrodynamics Simulations’, *PASA* **24**, 159–173.
- Price, D. J. (2012), ‘Smoothed particle hydrodynamics and magnetohydrodynamics’, *Journal of Computational Physics* **231**, 759–794.
- Pringle, J. E. (1981), ‘Accretion discs in astrophysics’, *Annual Reviews of Astronomy and Astrophysics* **19**, 137–162.
- Pringle, J. E. & Rees, M. J. (1972), ‘Accretion Disc Models for Compact X-Ray Sources’, *A&A* **21**, 1.

- Pritzl, B. J., Smith, H. A., Catelan, M. & Sweigart, A. V. (2001), 'Variable Stars in the Unusual, Metal-Rich, Globular Cluster NGC 6441', *AJ* **122**, 2600–2626.
- Robertson, H. P. (1937), 'Dynamical effects of radiation in the solar system', *MNRAS* **97**, 423.
- Rosswog, S. & Brüggem, M. (2007), *Introduction to High-Energy Astrophysics*, Cambridge University Press.
URL: <http://books.google.com.au/books?id=gyhmQgAACAAJ>
- Rothschild, R. E., Blanco, P. R., Gruber, D. E., Heindl, W. A., MacDonald, D. R., Marsden, D. C., Pelling, M. R., Wayne, L. R. & Hink, P. L. (1998), 'In-Flight Performance of the High-Energy X-Ray Timing Experiment on the Rossi X-Ray Timing Explorer', *ApJ* **496**, 538.
- Sakano, M., Koyama, K., Murakami, H., Maeda, Y. & Yamauchi, S. (2002), 'ASCA X-Ray Source Catalog in the Galactic Center Region', *ApJS* **138**, 19–34.
- Sala, G., Haberl, F., José, J., Parikh, A., Longland, R., Pardo, L. C. & Andersen, M. (2012), 'Constraints on the Mass and Radius of the Accreting Neutron Star in the Rapid Burster', *ApJ* **752**, 158.
- Sguera, V., Bazzano, A. & Bird, A. J. (2007), 'A long burst from SLX 1735-269 detected by INTEGRAL', *The Astronomer's Telegram* **1340**, 1.
- Shakura, N. I. & Sunyaev, R. A. (1973), 'Black holes in binary systems. Observational appearance.', *A&A* **24**, 337–355.
- Skinner, G. K., Foster, A. J., Willmore, A. P. & Eyles, C. J. (1990), 'Localization of one of the Galactic Centre X-ray burst sources', *MNRAS* **243**, 72–77.
- Smith, D. A., Morgan, E. H. & Bradt, H. (1997), 'Rossi X-Ray Timing Explorer Discovery of Coherent Millisecond Pulsations during an X-Ray Burst from KS 1731-260', *ApJ* **479**, L137.
- Spitkovsky, A., Levin, Y. & Ushomirsky, G. (2002), 'Propagation of Thermonuclear Flames on Rapidly Rotating Neutron Stars: Extreme Weather during Type I X-Ray Bursts', *ApJ* **566**, 1018–1038.
- Steiner, A. W., Lattimer, J. M. & Brown, E. F. (2010), 'The Equation of State from Observed Masses and Radii of Neutron Stars', *ApJ* **722**, 33–54.
- Stella, L. & Rosner, R. (1984), 'Magnetic field instabilities in accretion disks', *ApJ* **277**, 312–321.
- Strohmayer, T. & Bildsten, L. (2006), New views of thermonuclear bursts, *in* Lewin, W. H. G. and van der Klis, M., ed., 'Compact Stellar X-ray Sources', pp. 113–156.
- Strohmayer, T. E. & Brown, E. F. (2002), 'A Remarkable 3 Hour Thermonuclear Burst from 4U 1820-30', *ApJ* **566**, 1045–1059.
- Strohmayer, T. E., Jahoda, K., Giles, A. B. & Lee, U. (1997), 'Millisecond Pulsations from a Low-Mass X-Ray Binary in the Galactic Center Region', *ApJ* **486**, 355.
- Strohmayer, T. E. & Markwardt, C. B. (2002), 'Evidence for a Millisecond Pulsar in 4U 1636-53 during a Superburst', *ApJ* **577**, 337–345.
- Strohmayer, T. E., Zhang, W., Swank, J. H. & Lapidus, I. (1998), 'The Long-Term Stability of Oscillations during Thermonuclear X-Ray Bursts: Constraining the Binary X-Ray Mass Function', *ApJ* **503**, L147.

- Strohmayer, T. E., Zhang, W., Swank, J. H., Smale, A., Titarchuk, L., Day, C. & Lee, U. (1996), ‘Millisecond X-Ray Variability from an Accreting Neutron Star System’, *ApJ* **469**, L9.
- Suleimanov, V., Poutanen, J. & Werner, K. (2011), ‘X-ray bursting neutron star atmosphere models: spectra and color corrections’, *A&A* **527**, A139.
- Swank, J. H., Becker, R. H., Boldt, E. A., Holt, S. S., Pravdo, S. H. & Serlemitsos, P. J. (1977), ‘Spectral evolution of a long X-ray burst’, *ApJ* **212**, L73–L76.
- Swank, J. H., Becker, R. H., Pravdo, S. H., Saba, J. R. & Serlemitsos, P. J. (1976), ‘X-Ray Bursts’, *IAU Circ.* **3000**, 5.
- Sztajno, M., Fujimoto, M. Y., van Paradijs, J., Vacca, W. D., Lewin, W. H. G., Penninx, W. & Trumper, J. (1987), ‘Constraints on the mass-radius relation of the neutron star in 4U 1746-37/NGC 6441’, *MNRAS* **226**, 39–55.
- Taam, R. E. (1982), ‘Helium and combined hydrogen-helium shell flashes in the envelope of an accreting neutron star’, *Astrophysical Journal* **258**, 761–769.
- Tavani, M., Barbiellini, G., Argan, A., Bulgarelli, A., Caraveo, P., Chen, A., Cocco, V., Costa, E., de Paris, G., Del Monte, E., Di Cocco, G., Donnarumma, I., Feroci, M., Fiorini, M., Froyland, T., Fuschino, F., Galli, M., Gianotti, F., Giuliani, A., Evangelista, Y., Labanti, C., Lapshov, I., Lazzarotto, F., Lipari, P., Longo, F., Marisaldi, M., Mastropietro, M., Mauri, F., Mereghetti, S., Morelli, E., Morselli, A., Pacciani, L., Pellizzoni, A., Perotti, F., Picozza, P., Pontoni, C., Porrovecchio, G., Prest, M., Pucella, G., Rapisarda, M., Rossi, E., Rubini, A., Soffitta, P., Trifoglio, M., Trois, A., Vallazza, E., Vercellone, S., Zambra, A., Zanello, D., Giommi, P., Antonelli, A. & Pittori, C. (2008), ‘The AGILE space mission’, *Nuclear Instruments and Methods in Physics Research A* **588**, 52–62.
- Tawara, Y., Hirano, T., Kii, T., Matsuoka, M. & Murakami, T. (1984), ‘Long X-ray burst from GX 17 + 2’, *PASJ* **36**, 861–867.
- Taylor, M. B. (2005), TOPCAT STIL: Starlink Table/VOTable Processing Software, in P. Shopbell, M. Britton & R. Ebert, eds, ‘Astronomical Data Analysis Software and Systems XIV’, Vol. 347 of *Astronomical Society of the Pacific Conference Series*, p. 29.
- Taylor, M. B. (2006), STILTS - A Package for Command-Line Processing of Tabular Data, in C. Gabriel, C. Arviset, D. Ponz & S. Enrique, eds, ‘Astronomical Data Analysis Software and Systems XV’, Vol. 351 of *Astronomical Society of the Pacific Conference Series*, p. 666.
- The PARI Group (2011), *PARI/GP, version 2.5.0*, Bordeaux. available from <http://pari.math.u-bordeaux.fr/>.
- Thompson, M. J. (2006), *An introduction to astrophysical fluid dynamics*.
- Thompson, T. W. J., Rothschild, R. E., Tomsick, J. A. & Marshall, H. L. (2005), ‘Chandra and RXTE Spectra of the Burster GS 1826-238’, *ApJ* **634**, 1261–1271.
- Thorne, K. S. (1977), ‘The relativistic equations of stellar structure and evolution’, *Astrophysical Journal* **212**, 825–831.
- Titarchuk, L. (1994a), ‘Generalized Comptonization models and application to the recent high-energy observations’, *ApJ* **434**, 570–586.

- Titarchuk, L. (1994*b*), ‘On the spectra of X-ray bursters: Expansion and contraction stages’, *ApJ* **429**, 340–355.
- Treves, A., Maraschi, L. & Abramowicz, M. (1988), ‘Basic elements of the theory of accretion’, *PASP* **100**, 427–451.
- Vacca, W. D., Lewin, W. H. G. & van Paradijs, J. (1986), ‘Eddington luminosities and photospheric radius expansion during X-ray bursts from 4U/MXB 1820-30’, *MNRAS* **220**, 339–349.
- van Paradijs, J. (1979), ‘Possible observational constraints on the mass-radius relation of neutron stars’, *ApJ* **234**, 609–611.
- van Paradijs, J., Dotani, T., Tanaka, Y. & Tsuru, T. (1990), ‘A very energetic X-ray burst from 4U 2129 + 11 in M15’, *PASJ* **42**, 633–660.
- van Paradijs, J. & Lewin, H. G. (1986), ‘The interpretation of blackbody spectra and radii during X-ray bursts’, *A&A* **157**, L10–L12.
- Verbunt, F., Elson, R. & van Paradijs, J. (1984), ‘X-ray sources in globular clusters’, *MNRAS* **210**, 899–914.
- Walker, M. (1988), ‘Accretion-driven nonthermal phenomena’, *ApJ* **330**, L47–L49.
- Walker, M. A. (1992), ‘Radiation dynamics in X-ray binaries. I - Type 1 bursts. II - Type 2 bursts. III - Extremely compact objects’, *ApJ* **385**, 642–664.
- Walker, M. A. & Meszaros, P. (1989), ‘The dynamical influence of radiation in type 1 X-ray bursts’, *ApJ* **346**, 844–846.
- Wang, Z., Chakrabarty, D., Roche, P., Charles, P. A., Kuulkers, E., Shahbaz, T., Simpson, C., Forbes, D. A. & Helsdon, S. F. (2001), ‘The Optical Counterpart of the Accreting Millisecond Pulsar SAX J1808.4-3658 in Outburst: Constraints on the Binary Inclination’, *ApJ* **563**, L61–L64.
- Watts, A. L. (2012), ‘Thermonuclear Burst Oscillations’, *ARA&A* **50**, 609–640.
- Watts, A. L., Altamirano, D., Linares, M., Patruno, A., Casella, P., Cavecchi, Y., Degenaar, N., Rea, N., Soleri, P., van der Klis, M. & Wijnands, R. (2009), ‘Discovery of Burst Oscillations in the Intermittent Accretion-Powered Millisecond Pulsar HETE J1900.1-2455’, *ApJ* **698**, L174–L177.
- Weaver, T. A., Zimmerman, G. B. & Woosley, S. E. (1978), ‘Presupernova evolution of massive stars’, *ApJ* **225**, 1021–1029.
- Weinberg, N. N., Bildsten, L. & Schatz, H. (2006), ‘Exposing the Nuclear Burning Ashes of Radius Expansion Type I X-Ray Bursts’, *ApJ* **639**, 1018–1032.
- Werner, N., in’t Zand, J. J. M., Natalucci, L., Markwardt, C. B., Cornelisse, R., Bazzano, A., Cocchi, M., Heise, J. & Ubertini, P. (2004), ‘X-ray spectral evolution of SAX J1747.0-2853 during outburst activity and confirmation of its transient nature’, *A&A* **416**, 311–318.
- White, N. E. (1986), X-ray Spectral Formation in Low Mass X-ray Binaries, in K. O. Mason, M. G. Watson, & N. E. White, ed., ‘The Physics of Accretion onto Compact Objects’, Vol. 266 of *Lecture Notes in Physics*, Berlin Springer Verlag, p. 377.
- White, N. E. & Holt, S. S. (1982), ‘Accretion disk coronae’, *ApJ* **257**, 318–337.

- Wijnands, R., Strohmayer, T. & Franco, L. M. (2001), ‘Discovery of Nearly Coherent Oscillations with a Frequency of ~ 567 HZ during Type I X-Ray Bursts of the X-Ray Transient and Eclipsing Binary X1658-298’, *ApJ* **549**, L71–L75.
- Winkler, C., Courvoisier, T. J.-L., Di Cocco, G., Gehrels, N., Giménez, A., Grebenev, S., Hermsen, W., Mas-Hesse, J. M., Lebrun, F., Lund, N., Palumbo, G. G. C., Paul, J., Roques, J.-P., Schnopper, H., Schönfelder, V., Sunyaev, R., Teegarden, B., Ubertini, P., Vedrenne, G. & Dean, A. J. (2003), ‘The INTEGRAL mission’, *A&A* **411**, L1–L6.
- Worpel, H., Galloway, D. K. & Price, D. J. (2013), ‘Evidence for Accretion Rate Change during Type I X-Ray Bursts’, *ApJ* **772**, 94.
- Worpel, H., Galloway, D. K. & Price, D. J. (2015), ‘Evidence for enhanced persistent emission during sub-Eddington thermonuclear bursts’, *ArXiv e-prints* .
- Younes, G., Boirin, L. & Sabra, B. (2009), ‘An XMM-Newton view of the dipping low-mass X-ray binary XTE J1710-281’, *A&A* **502**, 905–912.
- Zamfir, M., Cumming, A. & Galloway, D. K. (2012), ‘Constraints on Neutron Star Mass and Radius in GS 1826-24 from Sub-Eddington X-Ray Bursts’, *ApJ* **749**, 69.
- Zhang, W., Jahoda, K., Kelley, R. L., Strohmayer, T. E., Swank, J. H. & Zhang, S. N. (1998), ‘Millisecond Oscillations in the Persistent and Bursting Flux of Aquila X-1 during an Outburst’, *ApJ* **495**, L9.